

Send documentation comments to n5kdocfeedback@cisco.com



New and Changed Information

This chapter provides release-specific information for each new and changed feature in the *Cisco Nexus 5000 Series NX-OS FCoE Operations Guide, Release 5.1(3)N1(1)*. The latest version of this document is available at the following Cisco website:

http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html

To check for additional information about Cisco NX-OS Release 5.x, see the *Cisco Nexus 5000 Series Switch NX-OS Release Notes* available at the following Cisco website:

http://www.cisco.com/en/US/products/ps9670/prod_release_notes_list.html

Table 1 lists the new and changed features for the *Cisco Nexus 5000 Series NX-OS FCoE Operations Guide, Release 5.1(3)N1(1)*.

Table 1 *New and Changed Features for Release 5.1(3)N1(1)*

Feature	Description	Changed in Release	Where Documented
FCoE NPV	This feature was introduced.	5.1(3)N1(1)	Chapter 6, “Using FCoE NPV”
Multi-Hop FCoE Solutions	Added a note for issues when port channeling multiple ports	5.1(3)N1(1)	Chapter 1, “Fibre Channel over Ethernet Operations”

Send documentation comments to n5kdocfeedback@cisco.com



CHAPTER 1

Fibre Channel over Ethernet Operations

This chapter includes the following sections:

- [Introduction, page 1-1](#)
- [FCoE Considerations, page 1-1](#)
- [FCoE Supported Topologies, page 1-14](#)
- [FCoE Operations, page 1-22](#)
- [Additional Information, page 1-28](#)

Introduction

The Cisco Nexus 5000 Series switch has supported FCoE since 2009. As the adoption of FCoE increases within the data center, there are design and operational considerations to take into effect. This document discusses these considerations and provides operational guidelines on how to deploy and implement an FCoE solution with Cisco Nexus 5000 Series switches.

FCoE Considerations

This section includes the following topics:

- [Preserving SAN Fabric Isolation, page 1-2](#)
- [FCoE and Spanning Tree Protocol Considerations, page 1-3](#)
- [FCoE and Virtual Port Channel \(vPC\) Considerations, page 1-5](#)
- [Changing Buffer Allocations for Longer Distance FCoE, page 1-8](#)
- [Consolidated Links And Dedicated Links for FCoE, page 1-9](#)
- [Cisco Nexus 5000 Series Switch FCoE Considerations, page 1-10](#)
- [Priority Flow Control and Enhanced Transmission Selection Considerations, page 1-12](#)
- [Cisco Nexus Interoperability, page 1-14](#)

Send documentation comments to n5kdocfeedback@cisco.com

Preserving SAN Fabric Isolation

High availability (HA) is a requirement in any data center design—whether it is accomplished through HA at the port level, supervisor level, or even at the physical network level. Fibre Channel Storage Area Networks (FC SANs) achieve high availability by building out two identical but physically separate networks commonly referred to as SAN A and SAN B (also called Fabric A and Fabric B). These networks, unlike Data Center LAN networks, are completely physically isolated from one another and have no knowledge of each other. Depending on host operating systems and drivers, traffic is able to be load balanced or “multi-pathed” between the two isolated networks, from the application side, in order to provide better service to the storage traffic. This required isolation is an important element in building FCoE networks along side the data center Ethernet LANs.

This section includes the following methods on preserving fabric isolation:

- [Maintaining Different FC-MAPs Per Fabric, page 1-2](#)
- [VLAN to VSAN Numbering, page 1-3](#)

Maintaining Different FC-MAPs Per Fabric



Note

We recommend using the “[VLAN to VSAN Numbering](#)” section on [page 1-3](#) method for preserving fabric isolation and leaving the FC-MAP default.

FC-MAP is a characteristic of a FCoE switch that identifies which fabric the switch belongs to. For instance, there can be an FC-MAP for Fabric A and a different FC-MAP for Fabric B. By configuring a specific FC-MAP value on a FCoE switch, it is possible to designate certain switches to belong to one fabric or another.

In order to maintain fabric isolation in an FCoE environment, it is recommended to use different FC-MAP values per SAN Fabric. Because the FC-MAP value of the Cisco Nexus 5000 Series switch is used in the addressing for FCoE-enabled devices, changing the FC-MAP value is a disruptive process to all hosts that are logged into the switch. Due to this disruption, it is recommended that the FC-MAP is configured as part of the initial switch set up.

By default, when the **feature fcoe** command is used to enable FCoE on a Cisco Nexus 5000 Series switch, a default FC-MAP is assigned to the switch. The simplest way to ensure SAN A and SAN B isolation between FCoE-enabled switches in the Ethernet fabric is to change the FC-MAP value to something other than the default for all switches belonging to Fabric B. This will prohibit FCoE switches from joining the wrong fabric and aide to providing the SAN isolation that is a requirement for FC and FCoE traffic.

To change the FC-MAP of a switch:

```
switch# configure terminal
switch(config)# fcoe fcmmap 0e.fc.2a
```



Note

Changing the FC-MAP value of a switch is disruptive to all attached FCoE hosts and it requires the hosts to login to the fabric again. Therefore, it is recommended to change the FC-MAP when the switch is installed and initially configured or during a maintenance window.



Note

The default value of the FC-MAP on a Cisco Nexus 5000 Series switch is 0E.FC.00. The configurable values for FC-MAP ranges from 0E.FC.00 to 0E.FC.FF.

Send documentation comments to n5kdocfeedback@cisco.com

VLAN to VSAN Numbering

When configuring an FCoE fabric, the first step is to create a VLAN to VSAN mapping which allows the FC traffic in a single VSAN to traverse the Ethernet network. It is a best practice to have dedicated VLANs for FCoE traffic in order to separate the storage traffic from all other Ethernet VLANs. It is also recommended not to assign VLAN 1, VSAN 1 or the configured native VLAN to the FCoE network. Typically those VLAN/VSANs are utilized for management traffic or for devices that have no other VLAN or VSAN assigned to them. Using VLAN 1 as an FCoE VLAN will not be supported on the Cisco Nexus 5000 Series switch running Cisco NX-OS release 5.0(1)N1(2) or a later release.

VLAN to VSAN mapping is a one-to-one relationship. Mapping multiple VSANs to a single VLAN instance is not supported. Note that both the VLAN instance and VSAN instance in an FCoE VLAN/VSAN mapping take up a hardware VLAN resource. Currently, there can be up to 31 VLAN/VSAN mappings supported on the Cisco Nexus 5000 Series switch. VLAN and VSAN numbering can range from 1-4096.

FCoE VLANs are different from typical Ethernet VLANs in that it acts more of a container for the storage traffic than anything else. MAC learning, broadcasts, or flooding do not occur and it does not map to a subnet. FCoE VLANs are simply used to carry the traffic for a specified FC VSAN and keep it separate from any other Ethernet VLANs that may be traversing the network.

In order to avoid confusion and service disruption in the event of a misconfiguration, it is recommended that you configure different FCoE VLAN and VSAN numbers for both SAN A and SAN B. Using the same VLAN or VSAN numbering between the two fabrics could result in the merging of both SAN fabrics in the event of a miss-configuration or miss-cabling. It is also best practice to only define SAN A VLANs on SAN A switches and vice-versa.

Host-facing FCoE ports must be configured as trunk ports carrying the native VLAN, FCoE VLAN and any other Ethernet VLANs necessary for the host application. These host facing ports should also be configured as spanning tree edge ports using the **spanning-tree port type edge [trunk]** interface-level command.



Note

- FCoE Initialization Protocol (FIP) uses the native VLAN and therefore all FCoE links should be trunked to carry the FCoE VLAN as well as the native VLAN.
- The FCoE VSAN must be configured and in the VSAN database of the Cisco Nexus 5000 Series switch prior to mapping it to a VLAN
- Enabling FCoE on VLAN 1 is NOT supported

FCoE and Spanning Tree Protocol Considerations

Native FC has no concept of a looped environment and therefore has no need for a protocol similar to the Spanning Tree Protocol (STP) in the Ethernet world. However, when placing FCoE onto an Ethernet fabric, STP is run on the FCoE VLANs connecting to a host (VF port) over a lossless Ethernet cloud. This lossless cloud could be made up of DCB bridges or FIP snooping devices. Because of this, there are certain recommendations for STP configurations that should be followed when deploying FCoE. The goal is to have isolated STP topologies between SAN A, SAN B, and the Ethernet fabric. This eliminates any Ethernet topology changes from affecting storage traffic.



Note

STP is not run on FCoE VLANs on VE port connections between two FCFs.

Send documentation comments to n5kdocfeedback@cisco.com



Note

Beginning with Cisco NXOS Release 5.0(1)N1(1) for the Cisco Nexus 5000 Series switch, STP is not run on FCoE VLANs on VF ports connecting directly to attached hosts (including host connections to a Cisco Nexus 2232 Fabric Extender). STP will continue to run on VF ports that connect to hosts through a DCB cloud or FIP snooping device.



Note

In Cisco NXOS Release 4.2(1)N2(1a) and earlier releases, STP runs on FCoE VLANs for any VF port connection (either direct attached hosts or hosts connected over a DCB cloud). Because of this, it is required to configure the VF port as a spanning-tree port type edge trunk

This section includes the following topics:

- [MST Instances For Dual Fabric FCoE Deployments, page 1-4](#)
- [PVST+ for Dual Fabric FCoE Deployments, page 1-4](#)

MST Instances For Dual Fabric FCoE Deployments

When running multi-instance STP in an Ethernet environment, it is required that all switches in the same MST region have the identical mapping of VLANs to instances. This does not require that all VLANs be defined on all switches. When running FCoE over an environment using MST, it is recommended to have a dedicated MST instances for the FCoE VLANs belonging to SAN A and a dedicated MST instance for the FCoE VLANs belonging to SAN B. These instances should be separate from any instances that include regular Ethernet VLANs. This example shows the FCoE VLANs in Fabric A are VLANs 20-29 and the FCoE VLANs in Fabric B are VLANs 30-39:

Spanning-tree MST configuration:

- name FCoE-Fabric
- revision 5
- instance 5 vlan 1-19,40-3967,4048-4093
- instance 10 vlan 20-29
- instance 15 vlan 30-39

In the above configuration, instance 5 maps to native Ethernet VLANs, instance 10 maps to the VLANs for Fabric A (20-29) and instance 15 maps to the VLANs for Fabric B (30-39).

Due to the MST configuration requirement, it will be necessary to have the same MST configuration, containing both the SAN A and SAN B instance, on all switches within the same MST region. This means that switches participating in SAN A will also contain an MST configuration with a separate instance for SAN B VLANs even though those SAN B VLANs will not be defined on the SAN A switch.

PVST+ for Dual Fabric FCoE Deployments

When running PVST, each VLAN already has its own spanning tree topology. Because FCoE traffic in each SAN fabric is defined by different individual VLANs, PVST+ will automatically isolate the spanning tree domains for the VLANs in SAN A, SAN B, as well as the Ethernet fabric.

Send documentation comments to n5kdocfeedback@cisco.com

FCoE and Virtual Port Channel (vPC) Considerations

Virtual Port Channeling (vPC) is an Ethernet feature that allows a single device to connect to multiple upstream devices and forward out all available links without the implications of spanning tree blocking paths due to Ethernet loops. vPC is useful in three situations:

1. Connecting a server to two upstream switches
2. Connecting a FEX to two upstream Nexus 5X00s
3. Connecting a switch to two upstream switches

The upstream switches in all scenarios must support the virtual port channel feature. The downstream device has no knowledge of the vPC and simply views the connection as a standard Ethernet port channel.

Though it is not possible to run FCoE traffic on top of a vPC because of the SAN A and SAN B physical isolation requirement in native FC, it is possible to run FCoE and vPC side-by-side on the same physical infrastructure from the host to the first-hop FCoE device. To configure this topology, the following must be considered:

- A host must connect to the upstream Cisco Nexus 5000 Series vPC pair switches using only 2 10G links – one attaching to a Cisco Nexus 5000 Series switch in Fabric A and one attaching to a Cisco Nexus 5000 Series switch in fabric B. This is commonly referred to a *single-port vPC* because only one port goes to each switch.
- Generation 2 CNAs are required in the host in order to support vPC topologies.



Note

- FCoE and vPC can run side-by-side only on single-port host-connected vPCs. FCoE and vPC's between a FEX and a Cisco Nexus 5000 Series switch or between two layers of switches is not supported.
- FCoE and vPCs containing more than one link to each access device is not supported. vPCs which coexist with FCoE must contain only a single link to each vPC peer device.
- vPC's across switches (FCFs) within the same SAN fabric is not supported. Each vPC peer must be part of different fabrics—one peer in SAN A and one peer in SAN B.

This section includes the following topics:

- [Required Teaming Drivers for vPC With CNAs, page 1-5](#)
- [Second Generation CNA Requirement, page 1-6](#)
- [View Of Ethernet Traffic And FC Traffic Through A CNA, page 1-6](#)
- [FCoE VLAN Configuration On A vPC, page 1-7](#)

Required Teaming Drivers for vPC With CNAs

When connecting a host to an upstream vPC switch pair, the only requirement from the host side is to support link aggregation on the NIC interfaces. This can be accomplished using link aggregation control protocol (LACP) or standard 802.3ad *port channel mode on* behavior. It is important to check that either the host operating system or the native CNA hardware supports one of these options.

Send documentation comments to n5kdocfeedback@cisco.com

Second Generation CNA Requirement

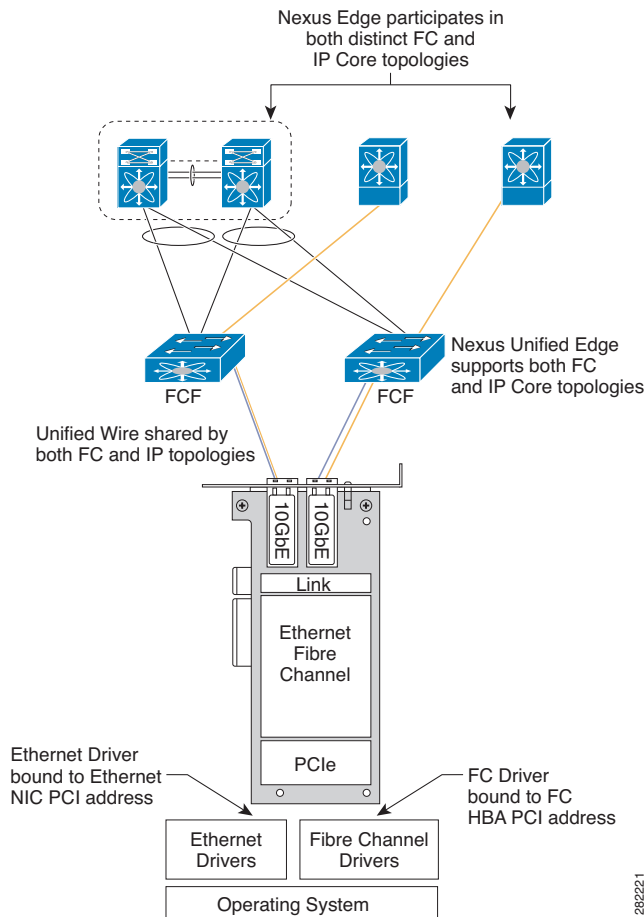
When connecting a host containing a CNA to upstream Cisco Nexus 5000 Series switches configured in a vPC, 2nd generation CNAs are required from both Emulex and QLogic. This is regardless of the presence of FCoE traffic on the host connections. These 2nd generation CNAs are also required when connecting to a Cisco Nexus 2232 Fabric Extender with a vPC (Ethernet only), FCoE, or FCoE+vPC configuration from a host connection.

View Of Ethernet Traffic And FC Traffic Through A CNA

Currently CNAs present two different types of adapters to the host operating system: Ethernet NICs and Fibre Channel HBAs. Though these adapters physically correspond to the same 10GE port on a CNA, to the operating system, it will appear as two completely separate and physically isolated interfaces. Because of this adapter port virtualization, it is possible to build two separate topologies based on traffic type: one for the Ethernet fabric using the NICs and one for the FC fabric using the HBAs.

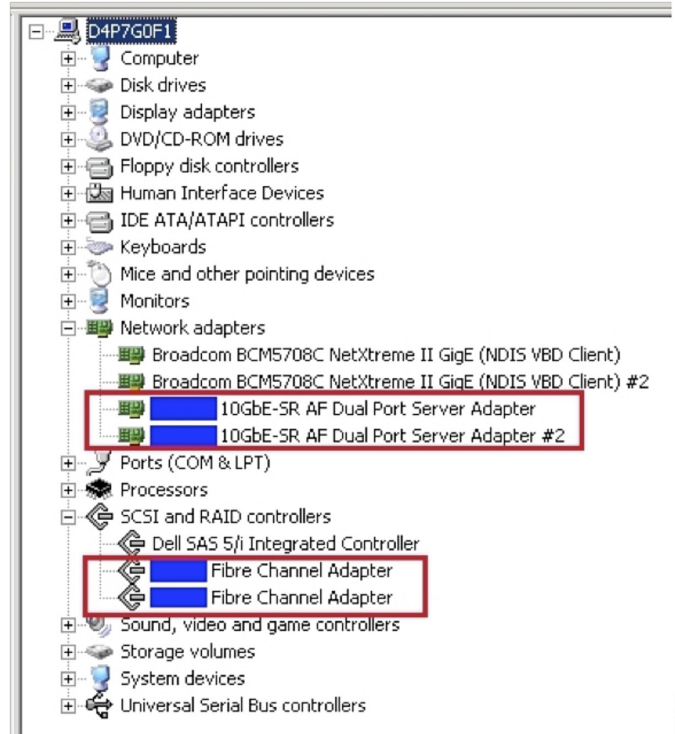
For FCoE and vPC to run side-by-side from the host, the port channel would be configured on the NICs interfaces presented and SAN multi-pathing or other SAN HA mechanisms would be configured on the FC HBAs presented to the OS by the CNA. Today, it is required that only 2X10GE links be used in a host side vPC port channel when running FCoE on the same wires. Each 10GE link will be used to provide a single connection to each upstream vPC switch.

Figure 1-1 Ethernet And FC Traffic Through A CNA



Send documentation comments to n5kdocfeedback@cisco.com

Figure 1-2 Adapter Control Panel Display



Note


- VPC + FCoE over a consolidated wire from the host requires the host supports port channels capabilities (LACP or “port channel mode ON”). Please check with specific CNA and OS vendors for a support matrix.
- VPC + FCoE over a consolidated wire are only supported between a host and either the first hop Nexus 5000 or Nexus 5000/2232 pair. VPC and FCoE on a consolidated wire is NOT supported beyond the access layer or when connecting a host to the Nexus 7000 platform.
- vPC and FCoE can not coexist on the same wire beyond any first hop access device.

FCoE VLAN Configuration On A vPC

Typically, interfaces belonging to the same port channel are required to have the same port configuration. This includes VLAN configuration. However, in order to maintain fabric separation alongside vPC connections, it is necessary to declare the FCoE VLAN for SAN A on one uplink and the FCoE VLAN for SAN B on the other uplink. This is a recommended best practice configuration. [Figure 1-3](#) shows the hosts connected to a Cisco Nexus 5000 Series switch running vPC and FCoE simultaneously.

Send documentation comments to n5kdocfeedback@cisco.com

Figure 1-3 FCoE VLAN Configuration In A vPC Topology



Changing Buffer Allocations for Longer Distance FCoE

Beginning with the Cisco NXOS Release 5.0(1)N1(1) for the Cisco Nexus 5000 Series switch, it is possible to tune the port buffer allocation and xon and xoff thresholds in order to support increased distance between VE ports. The default distance configured for each port when configured to carry FCoE traffic (or any “no-drop” traffic) is 300 meters. This supported distance is based on the amount of available buffer space allocated to catch frames in flight between the time a PAUSE is initiated towards a downstream device and the time that downstream devices processes that PAUSE frame and stops sending frames. This per port buffer allocation and configuration must match between the two ports on either end of the link (including host CNA ports as well). This is similar to the way buffer-to-buffer credits is initialized between two devices in a native FC environment.

The current xon threshold and buffer size allocated for FCoE is such that $\text{buffer-size} - \text{xon} = \sim 300$ meters worth of FCoE frames. The default configuration parameters for the class-fcoe (or any no-drop class) on the Nexus 5000 series switch is shown below:

- qos-group 1
- q-size: 76800, HW MTU: 2400 (2240 configured)
- drop-type: no-drop, xon: 128, xoff: 240

In order to support a distance of 3000m for FCoE traffic between two FCoE capable switches (connecting two FCFs with VE ports), the buffer allocation as well as the xon and xoff values need to be altered for the FCoE class of services: class-fcoe. This can be accomplished by editing the quality of service configuration. An example of this configuration can be found in the “Configuring NO-Drop Buffer Threshold” section of the Nexus 5000 Configuration Guide:

http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/qos/502_n1_1/Cisco_Nexus_5000_Series_NX-OS_Quality_of_Service_Configuration_Guide_Rel_502_N1_1.pdf

The necessary thresholds for support no-drop service up to 3000m is outlined in the table below:

Send documentation comments to n5kdocfeedback@cisco.com

Configuration for 3000m no-drop class	Buffer size	Pause Threshold (XOFF)	Resume Threshold (XON)
Nexus 5000 Series	143680 bytes	58860 bytes	38400 bytes
Nexus 5500 Platform	152000 bytes	103360 bytes	83520 bytes

Consolidated Links And Dedicated Links for FCoE

Because FCoE uses the Ethernet fabric for transport, there is the possibility of consolidating both Ethernet LAN traffic and Storage SAN traffic onto the same infrastructure. There are multiple levels of consolidation; wire consolidation and device consolidation are two of the most common and are discussed below.


Link Consolidation refers to when Ethernet LAN traffic and Storage SAN traffic are sharing the same physical wire between host and switch or between two switches.

Device consolidation refers to when Ethernet LAN traffic and Storage SAN traffic are passing through the same switching device but maintain isolation through the use of dedicated wires or switch ports.

The topologies discussed throughout this guide will mention two terms to describe the scope of FCoE traffic: consolidated link – where FCoE and native Ethernet traffic simultaneously use the same link -- and dedicated link – where FCoE and native Ethernet traffic use two separate DCB Ethernet links. The following sections will discuss the different places in the Data Center Network where consolidated and dedicated links make sense.

Figure 1-4 shows an example of consolidated vs dedicated links. The wires running from the host to the access devices are consolidated links carrying both Ethernet and FCoE traffic. Moving from the access to aggregation, there are dedicated links: blue wires dedicated to the Ethernet traffic and orange wires dedicated to FCoE traffic only.

Figure 1-4 Consolidated And Dedicated Links



This section includes the following topics:

Send documentation comments to n5kdocfeedback@cisco.com

- [Where Consolidated Links Makes Sense, page 1-10](#)
- [Where Dedicated Wires Makes Sense, page 1-10](#)

Where Consolidated Links Makes Sense

One of the benefits of FCoE at the access layer is the ability to consolidate the FC SAN and Ethernet LAN onto the same physical wires and same physical devices. This consolidation lends to a large CapEx savings by reducing the number of access switches, host adapters, cables and optics required to run both LAN and SAN networks within the data center. This consolidation is made possible due to the excess bandwidth that 10GE to the server is able to provide. Because very few servers in the Data Center today are pushing 10-Gigabit Ethernet of Ethernet-only traffic, there is room for the added storage traffic to share these common wires without impacting the performance of the host application.

Also, due to the CNA behavior and ability to present to the host application different physical devices corresponding to both LAN and SAN networks, it is possible to separate Ethernet HA from FC HA at the host level. This is accomplished by being able to use separate Ethernet teaming options on the NICs while using separate FC multi-pathing options on the HBAs. Depending on the operating system and CNA being used, these teaming options will vary.

Moving beyond the access layer, oversubscription ratios and Ethernet bandwidth provisioning will determine the amount of excess bandwidth available and the benefit of running consolidated links vs. dedicated links within the Data Center.

Where Dedicated Wires Makes Sense

High Availability requirements in LAN and SAN networks differ considerably. Where in Ethernet, HA is achieved by multi-homing devices to one another (partial/full Mesh), in Fibre Channel (and FCoE), HA is achieved by building two physically isolated networks. Both of these requirements must be met in a network that combines FCoE and Ethernet.

There have been multiple enhancements to the Ethernet HA model that improves on Ethernet Data Center design by overcoming some of the challenges of the Spanning Tree protocol. One example of this is the virtual Port Channeling feature found in the Nexus product suite. The nature of vPC is to be able to forward out multiple paths to multiple upstream devices without spanning tree blocking any of the uplinks. While this is great for Ethernet traffic, it breaks the SAN A/SAN B isolation required for FC/FCoE.

Therefore, it is often beneficial to use dedicated wires for Ethernet traffic and Storage traffic independently. With dedicated wires, the Ethernet links can be configured to take advantage of advanced Ethernet features such as vPC and the storage links can be configured based on the fabric isolation requirement. This is especially common when connected access switches to upstream LAN aggregation/SAN core devices.

Cisco Nexus 5000 Series Switch FCoE Considerations

The Cisco Nexus 5000 Series switches include a Unified Port Controller (UPC) ASIC responsible for the handling the forwarding decisions and buffering for multiple 10-Gigabit Ethernet ports:

- The Cisco Nexus 5000 Platform switches (the Cisco Nexus 5010 switch and the Cisco Nexus 5020 switch) include the first generation UPC ASIC.
- The Cisco Nexus 5500 Platform switches (the Cisco Nexus 5548P switch, Nexus 5548UP switch, and Nexus 5596UP switch) include the second generation UPC ASIC.

Send documentation comments to n5kdocfeedback@cisco.com

The following sections discuss the differences between the first and second generation architectures that relate to FCoE configuration and supported topologies.

This section includes the following topics:

- [VLAN Scalability, page 1-11](#)
- [FCoE QoS Configuration, page 1-11](#)
- [Unified Port Options, page 1-11](#)

VLAN Scalability

One of the differences between the first and second generation ASICs is the number of available VLAN resources available. The first generation ASICs support up to 512 VLANs (507 of which are user configurable). With the second generation ASIC, the available VLAN number has increased from 512 to 4096. Currently, 31 VLANs and 31 VSANs are supported for FCoE VLAN/VSAN mappings on both generations.



Note

The VLAN and the VSAN in an FCoE VLAN/VSAN mapping consume a hardware VLAN resources.

FCoE QoS Configuration

The Nexus 5000 Series switches always reserve some buffer space for FCoE traffic. When you enable the FCoE feature on Nexus 5000 Series switch, Nexus automatically configures the necessary QoS policy and buffer allocations using the reserved buffers.

The Nexus 5500 Series switches allow all available port buffers to be configured based on traffic needs. This allows you to create a custom FCoE policy that can use any available buffers.

When you enable FCoE on a Nexus 5500 Series switch, the system looks for a custom QoS policy. If it does not find one, it automatically uses the default QoS configuration shown below:

```
switch(config-sys-qos) # service-policy type qos input fcoe-default-in-policy
switch(config-sys-qos) # service-policy type queuing input fcoe-default-in-policy
switch(config-sys-qos) # service-policy type queuing output fcoe-default-out-policy
switch(config-sys-qos) # service-policy type network-qos fcoe-default-nq-policy
```

For more information, see the Cisco Nexus 5000 Series NX-OS Quality of Service Configuration Guide, which is available from:

http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html

Unified Port Options

Unified ports are capable of operating a 1- and 10-Gigabit Ethernet or 1-, 2-, and 4-Gigabit or 2-, 4-, and 8-Gigabit FC (depending on the transceiver used) which provide more configuration flexibility. Unified ports no longer require you to purchase a set number of FC ports through an expansion module. Unified Ports are available in the expansion module on the Cisco Nexus 5548P switch and the Nexus 5548UP platform as well as on all base ports of the Cisco Nexus 5596UP switch. There are configuration requirements that must be carefully followed when utilizing unified ports.

Ports of a similar type, either Ethernet or FC, must be configured in a contiguous sequence. Changes to the port-type require a switch reboot or expansion module reboot depending on where the unified ports are configured. For this reason, careful planning should be done when first configuring the switch. Cisco

Send documentation comments to n5kdocfeedback@cisco.com

recommends as a best practice to start Ethernet port configurations at one end of the platform (from Eth1/1 counting up) and the necessary Fibre Channel ports configured from the opposite end of the platform (Eth 1/48 counting down).

For additional information on configuring unified ports, see the *Unified Port Configurations on Cisco Nexus 5500 Platform Switches* documentation.

Priority Flow Control and Enhanced Transmission Selection Considerations

Both Priority Flow Control (PFC) and Enhanced Transmission Selection (ETS) are part of the IEEE 802.1Q Enhance Ethernet Standards that are currently in the final stages of standardization. Both PFC and ETS are support on all Cisco Nexus 5000 Series switches. PFC is class of service (COS) based PAUSE allowing for FCoE traffic assigned to a specific COS value to retain the lossless qualities which are required for the FC protocol. ETS is a mechanism for dividing a 10-Gigabit Ethernet link into multiple lanes based on the COS value and allocating the necessary bandwidth requirements which are honored in the presence of congestion. ETS prevents situations where default traffic would interfere with higher priority traffic.

PFC and ETS are often used in today's FCoE networks to provide lossless transport and dedicated bandwidth for FCoE traffic. However, they are not specific to FCoE and have many uses outside of an FCoE environment for providing specific levels of service to specific traffic classes.

This section includes the following topics:

- [Default PFC and ETS Settings, page 1-12](#)
- [Changing PFC and ETS Settings, page 1-12](#)
- [Host-Side Considerations For Altering PFC And ETS Settings, page 1-13](#)

Default PFC and ETS Settings

PFC and ETS both use the Class of Service (COS) bits in order to classify between traffic types. There are 8 COS values in the IEEE 802.1Q standard trunking header for Ethernet frames. The Cisco Nexus 5000 Series switch allows you to manually configure 6 classes. Up to 4 of the 6 user configurable classes can be designated as no-drop classes of service, meaning that in the event of port congestions, traffic belonging to the no-drop classes will pause to prohibit packet drop.

By default, the Nexus 5000 Platform as well as other vendor's FCoE products have decided on COS value of 3 for FCoE traffic. When FCoE is enabled on the Cisco Nexus 5000 Series switch, COS 3 is automatically configured for no-drop service (PFC setting) as well as a guarantee of 50% of the bandwidth in the case of congestion (ETS setting). It is best practice to leave the default COS value of 3 for FCoE traffic due to the agreement between vendors to support this as a "no-drop" class.

In the event that other traffic already exists within the network that is using the COS value of 3 or there is another reason to move FCoE traffic from COS 3, this can be changed through a Quality of Service configuration.

Changing PFC and ETS Settings

PFC and ETS settings are configured and changed in the Quality of Service configuration on the Nexus 5000 Series switch. This example shows a QoS configuration that changes the FCoE no-drop class of service to COS 4 as the reserved bandwidth for FCoE to 20% of the 10-Gigabit Ethernet link:

Send documentation comments to n5kdocfeedback@cisco.com

Step 1 Create classification rules first by defining and applying policy-map type qos:

```
N5k(config)# class-map type qos class-lossless
N5k(config-cmap-qos)# match cos 4
N5k(config-cmap-qos)# policy-map type qos policy-lossless
N5k(config-pmap-qos)# class type qos class-lossless
N5k(config-pmap-c-qos)# set qos-group 7
N5k(config-pmap-uf)# system qos
N5k(config-sys-qos)# service-policy type qos input policy-lossless
```

Step 2 Define and apply policy-map type network:

```
N5k(config-pmap-qos)# class type network-qos policy-lossless
N5k(config-cmap-uf)# match qos-group 7
N5k(config-cmap-uf)# policy-map type network-qos policy-lossless
N5k(config-pmap-uf)# class type network-qos class-lossless
N5k(config-pmap-uf-c)# pause no-drop
N5k(config-pmap-uf)# system qos
N5k(config-sys-qos)# service-policy type network-qos policy-lossless
```

Step 3 Create classification rules first by defining and applying policy-map type qos:

```
N5k(config)# class-map type queuing class-voice
N5k(config-cmap-que)# match qos-group 2
N5k(config-cmap-que)# class-map type queuing class-high
N5k(config-cmap-que)# match qos-group 3
N5k(config-cmap-que)# class-map type queuing class-low
N5k(config-cmap-que)# match qos-group 7
N5k(config-cmap-que)# exit
```

Step 4 Create classification rules for the individual classes:

```
N5k(config)# policy-map type queuing policy-BW
N5k(config-pmap-que)# class type queuing class-voice
N5k(config-pmap-c-que)# priority
N5k(config-pmap-c-que)# class type queuing class-voice
N5k(config-pmap-c-que)# bandwidth percent 20
N5k(config-pmap-c-que)# class type queuing class-high
N5k(config-pmap-c-que)# bandwidth percent 40
N5k(config-pmap-c-que)# class type queuing class-low
N5k(config-pmap-c-que)# bandwidth percent 10
N5k(config-pmap-c-que)# class type queuing class-fcoe
N5k(config-pmap-c-que)# bandwidth percent 30
N5k(config-pmap-c-que)# class type queuing class-default
N5k(config-pmap-c-que)# bandwidth percent 0
N5k(config-pmap-c-que)# system qos
N5k(config-sys-qos)# service-policy type queuing output policy-BW
```

Host-Side Considerations For Altering PFC And ETS Settings

Data Center Bridging eXchange (DCBX) protocol is another portion of the IEEE 802.1Q Data Center Bridging (DCB) standard currently in review by the Ethernet standards body. DCBX is a protocol that runs between DCB-capable devices to ensure that PFC and ETS settings are configured consistently between DCB peers. DCB can also be used as a way to configure DCB peer devices from a central switching location. CNAs that support DCB-*willing* are configured to accept the DCB configurations (including PFC and ETS settings) of the upstream DCB switching device. This greatly simplifies management and configuration of DCB and FCoE devices.

Send documentation comments to n5kdocfeedback@cisco.com

If changing the default configuration for FCoE traffic on the Cisco Nexus 5000 Series switch, it is possible for the switch to relay these configuration changes to any connected CNAs using the DCBX protocol. It is necessary that the CNA vendor and platform support DCBX in a *willing* mode in order for this to take place. Please check with the individual CNA vendors on whether they support receiving DCBX configurations for a network device.

If the CNA does not support a method of DCB-willing, in order to change from a default PFC and ETS configuration, it is required to manually alter the configuration of the Nexus 5000 Series as well as the downstream CNA device so that they are the same. Depending on the CNA, different tools or commands will be used to change these settings.



Note

If the DCBX negotiation fails between a host and switch or between a switch and switch, the PFC setting will not be set on the Nexus 5000 Series switch and the vFC interfaces will remain down until the DCB configuration matches.



Note

Though the DCBX standard states that there are 8 possible no-drop lanes, CNA vendors differ on the number of COS values that are supported for FCoE and no-drop service today. Check with the CNA vendor for the correct number of supported FCoE and no-drop classes.

Cisco Nexus Interoperability

For information on interoperability, see the [Cisco Data Center Interoperability Support Matrix](#).

FCoE Supported Topologies

This section includes the following topics:

- [Single-Hop FCoE Deployment Topologies](#), page 1-14
- [Multi-Hop FCoE Solutions](#), page 1-21

Single-Hop FCoE Deployment Topologies

There are two possible single-hop solutions when deploying FCoE with a Cisco Nexus 5000 Series switch and Cisco Nexus 2000 Series Fabric Extender. The first solution is referred to as “direct connect” where a host is directly connected to the first hop converged access switch. The second single hop solution deploys a FEX between the server and the first hop switch. Because the FEX acts as a remote line card to the parent switch and has no local switching capabilities, it is not considered a hop in the Ethernet or Storage topologies. The following section outlines in detail the current single hop deployment options and configurations which are supported with the switch and FEX today.

This section includes the following topics:

- [Switch Mode and NPV Mode](#), page 1-15
- [vPC and Active/Standby](#), page 1-16
- [Direct Attached CNAs With Active/Standby Ethernet Topologies](#), page 1-16
- [Direct Attached CNAs With vPC Ethernet Topologies](#), page 1-17

Send documentation comments to n5kdocfeedback@cisco.com

- [Cisco Nexus 5000 Series Switch and Cisco Nexus 2000 Fabric Extender Topologies, page 1-17](#)
- [FIP Snooping Bridges, page 1-18](#)
- [Cisco Nexus 4000 Series Switch To Cisco Nexus 5000 Series Switch FCoE With Consolidated Links, page 1-19](#)
- [Cisco Nexus 4000 Series Switch Connected To A Cisco Nexus 5000 Series Switch FCoE With Dedicated Wires, page 1-20](#)

Switch Mode and NPV Mode

The Cisco Nexus 5000 Series switch has two modes of operation relating to storage traffic forwarding: switch mode and N-Port Virtualizer (NPV) mode. This is the same as the modes of operation available on the Cisco Multiprotocol Director Series (MDS) Fibre Channel switches. The default mode on both platforms is “switch” mode. In the following topologies, the Cisco Nexus 5000 Series switch can either be in switch or NPV mode. The only requirement for a Cisco Nexus 5000 Series switch in NPV mode is that the upstream device supports the standard N-Port ID Virtualization (NPIV) functionality.

When the Cisco Nexus 5000 Series switch is operating in switch mode, all fabric services, for example, FSPF, zoning or DNS, are native on the access device. This means that all forwarding decisions are made by FSPF running on the switch. This mode also means that the switch consumes a Domain ID within the Fibre Channel Fabric. Limitations exist as to the number of Domain IDs that are supported within a single fabric. Specific domain ID limitations are defined by the storage vendors and OSM partners.

NPV defines the ability for a Fibre Channel switch to act as a proxy for both FLOGIs and forwarding decision and pass those duties to an upstream device. This upstream device must be capable of running NPIV which is an FC standard allowing multiple FCiDs to be handed out a single FC port. The benefit of an NPV device in a FC network is the elimination of the domain ID and therefore the ability to add more FC switches to a fabric without exceeding the supported Domain ID limitation.

The Cisco Nexus 5000 Series switch can also operate in NPV mode. When NPV is enabled on the switch, no FC fabric services are run locally on the platform and instead, forwarding and zoning services are handled by the upstream NPIV device. To avoid interoperability challenges when connecting a switch to a non-Cisco SAN core switch, Cisco recommends that the switch be configured in NPV mode.

Enabling NPV on the switch is a disruptive process and should be done at the time of initial set up to avoid any disruption to the fabric. Because enabling NPV requires a switch reboot and erases the current running configuration, be sure to save the current running configuration to an external text file so that it can be reapplied after the reboot occurs if enabling NPV after the initial set up of the switch.

Changing between switch mode and NPV mode can be done using the following commands:

To enable NPV mode:

```
switch# feature npv
```

To disable NPV mode (return to switch mode):

```
switch# no feature npv
```



Note

Running NPV on the switch requires that the upstream connected device has NPIV functionality enabled



Note

FC or FCoE hosts conversing with an FC or FCoE storage devices connected to the same switch in NPV is NOT supported.

Send documentation comments to n5kdocfeedback@cisco.com

vPC and Active/Standby

Host facing interfaces on the Nexus 5000 Series switch can provide connections to servers in a couple of different ways: single attached NICs for single attached hosts, active-standby NIC teaming for dual-homed servers and vPC for dual-homed servers. This guide focuses on the dual-homed server options as FC requires two independent paths to storage: Fabric A and Fabric B.


Active/Standby connections refer to servers that are dual-homed to an Ethernet LAN but only actively forwarding out one link. The second link is used as back-up in case of a failure but does not actively forward traffic unless a failure occurs. vPC is a technology introduced by Cisco Nexus products that allows a dual homed server to actively forward out both Ethernet links simultaneously. The benefits of vPC is that it gives servers access to twice as much bandwidth as in an active/standby configuration and also has the ability to converge faster than Spanning-tree in the event of a failure.

Based on the Ethernet high availability requirement, LAN admins may choose to attached servers using active/standby connections or vPC connections. Regardless of the method use to dual home a server, FCoE can co-exist with both of these topologies.

Direct Attached CNAs With Active/Standby Ethernet Topologies

Figure 1-5 shows a topology where a dual-port CNA is connecting to two switches in an active/standby configuration. Although Ethernet traffic will only traverse one link in this configuration, the FCoE traffic will be forwarded out both paths to the fabric. This is because of the way the CNA is able to differentiate between the NIC adapters for Ethernet and FC adapters for FC/FCoE. For more information on the CNA view of the Ethernet NICs and storage HBAs, see the [“View Of Ethernet Traffic And FC Traffic Through A CNA”](#) section on page 1-6.

Figure 1-5 *Dual-Port CNA Connecting To Two Cisco Nexus 5000 Series Switches In An Active/Standby Topology*




Send documentation comments to n5kdocfeedback@cisco.com

Direct Attached CNAs With vPC Ethernet Topologies

Figure 1-6 shows a topology where a dual-port CNA is connecting to two switches in a vPC configuration where only a single port connects the CNA to each switch. The operating system is able to see the Ethernet aspects of these two physical ports and port channel the Ethernet traffic coming out of the server. The FC traffic is still mapped to each link separately – one 10-Gigabit link transporting Fabric A traffic and the other 10-Gigabit link transporting Fabric B traffic. For more information on the CNA view of the Ethernet NICs and Storage HBAs, see the “View Of Ethernet Traffic And FC Traffic Through A CNA” section on page 1-6.

Figure 1-6 *Dual-Port CNA Connecting To Two Cisco Nexus 5000 Series Switches In A vPC Topology*



Note

Direct-connect FCoE (a CNA that is directly connected to a Cisco Nexus 5000 Series switch switchport) is not supported on a port channel interface configured to have more than one member port. Directly connected FCoE devices are supported over virtual port channels where a single link from each CNA port connects through to each upstream switch or fabric extender.

Cisco Nexus 5000 Series Switch and Cisco Nexus 2000 Fabric Extender Topologies


The Nexus 2232 Fabric Extender acts as a remote line card to the parent Cisco Nexus 5000 Series switch. The Nexus 2232 Fabric Extender has 32 10-Gigabit Ethernet host facing interfaces, all of which support lossless Ethernet and FCoE. Supporting FCoE over a Cisco Nexus 5000 Series switch and FEX topology has the following requirements:

- Each Nexus 2232 Fabric Extender running FCoE must be single-homed to the upstream parent switch.
- Generation 2 (FIP Enabled) CNAs are required for host connections to the Cisco Nexus 2232 Fabric Extender host interfaces.

Send documentation comments to n5kdocfeedback@cisco.com

Adding the Cisco Nexus 2232 Fabric Extender into the FCoE topology does not change the supported configurations. Hosts can be connected to the Cisco Nexus 2232 Fabric Extender using active/standby Ethernet connections or over vPC connections. Figure 1-7 shows the supported topology.

Figure 1-7 Hosts Connected To The Cisco Nexus 2232 Fabric Extender Using Active/Standby Ethernet Connections or vPC Connections



Note

FCoE is not supported on a FEX interface or port channel interfaces when the FEX is connected to two switches in a FEX active-active topology.

FIP Snooping Bridges

FIP Snooping Bridges (FSBs) are lossless Ethernet bridges that are capable of watching a FIP conversation between a CNA and FCF. They have no FC/FCoE forwarding logic capabilities but instead “snoop” FIP packets and watch the FIP conversation, including FLOGI/LOGIN, between the CNA and FCF. Once a FIP snooping bridge sees a CNA login to the FC/FCoE fabric through a specific FCF, it dynamically creates an access list to guarantee that the communication between that CNA and FCF will remain point-to-point. FIP snooping is a security precaution used when transversing lossless Ethernet bridges to ensure that rogue devices can not enter the data center network and pretend to be an FCF.

It is important to note that FSBs are Layer 2 Lossless Ethernet bridges that have been enhanced to dynamically create ACLs based on the FIP communication that is seen within the fabric. FSBs have no knowledge of FC/FCoE protocols or services and do not forward FCoE traffic based on FSPF. Instead, all traffic runs over the Layer 2 protocol (STP) and is switched based on MAC address.

Send documentation comments to n5kdocfeedback@cisco.com

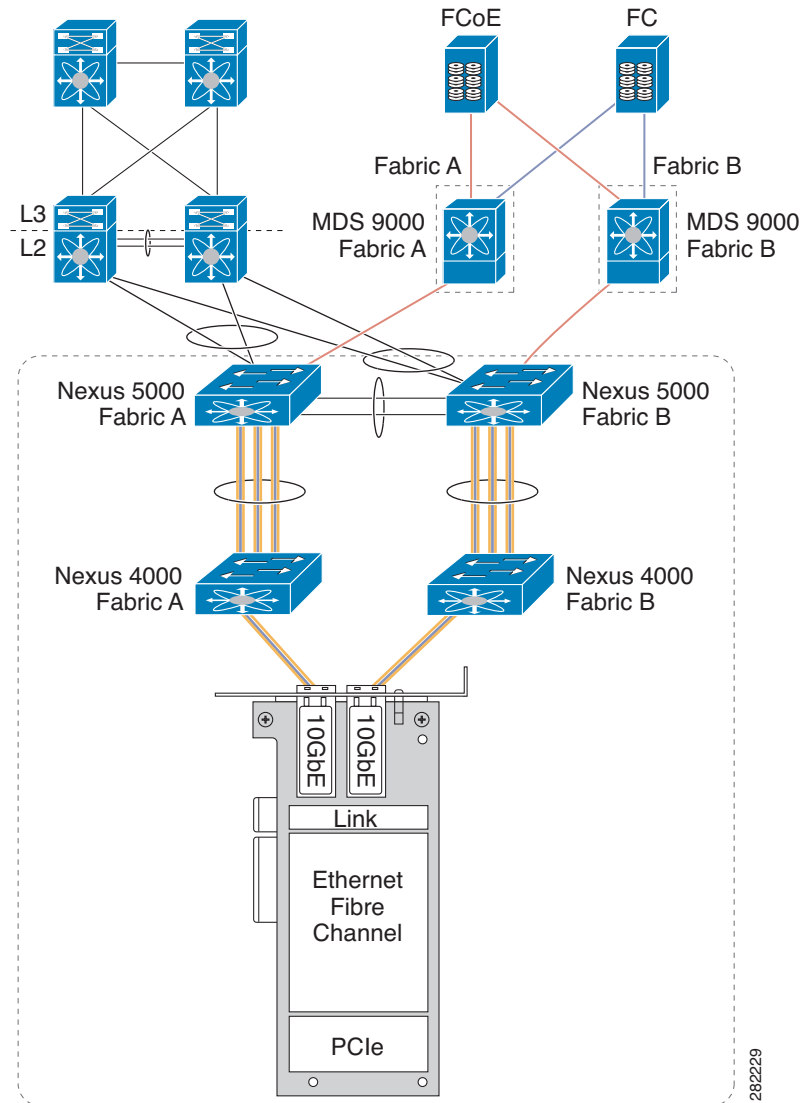
The Cisco Nexus 4000 Series switch is a FIP Snooping device for IBM blade chassis and must be connected to a Cisco Nexus 5000 Series FCF switch in order to support passing FCoE frames. The Cisco Nexus 4000 Series switch has 14 down-facing 10-Gigabit ports connecting to each of the 14 blade servers and 6 10-Gigabit Ethernet uplink ports used to connect to a Cisco Nexus 5000 Series switch. [Figure 1-8](#) and [Figure 1-9](#) shows the two supported configurations when connecting a Cisco Nexus 4000 Series switch FIP Snooping bridge to a Cisco Nexus 5000 Series FCF switch:

Cisco Nexus 4000 Series Switch To Cisco Nexus 5000 Series Switch FCoE With Consolidated Links

[Figure 1-8](#) shows a Cisco Nexus 4000 Series switch connected to a Cisco Nexus 5000 Series switch using consolidated links where both FCoE and Ethernet traffic are utilizing the same link simultaneously. Because FCoE requires fabric separation, the Ethernet traffic must also only follow one path and can not take advantage of other Ethernet HA technologies such as vPC.

Send documentation comments to n5kdocfeedback@cisco.com

Figure 1-8 Cisco Nexus 4000 Series Switch Connected To A Cisco Nexus 5000 Series Switch FCoE With Consolidated Links

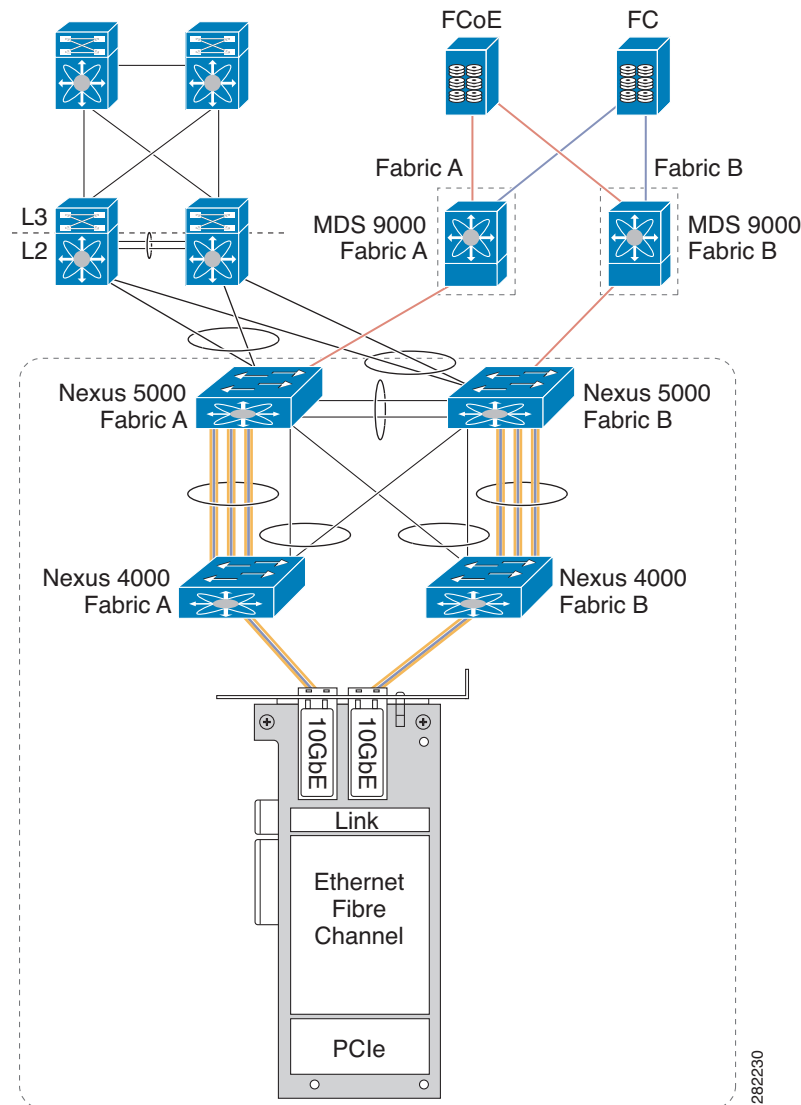


Cisco Nexus 4000 Series Switch Connected To A Cisco Nexus 5000 Series Switch FCoE With Dedicated Wires

Figure 1-9 shows the Cisco Nexus 4000 Series switches connecting to Cisco Nexus 5000 Series switches using dedicated links; blue links are Ethernet ONLY links and pink and blue links are FCoE-only links. There are no consolidated links shown in Figure 1-9. The benefit of running dedicated links between the Cisco Nexus 4000 Series switches and Cisco Nexus 5000 Series switches in this topology is the fact that both storage and Ethernet traffic are able to take advantage of their respective HA models. Ethernet traffic is multi-homed to the upstream switches and using vPC to forward out all available paths while FCoE is maintaining fabric isolation through the Ethernet network.

Send documentation comments to n5kdocfeedback@cisco.com

Figure 1-9 Cisco Nexus 4000 Series Switch Connected To A Cisco Nexus 5000 Series Switch FCoE With Dedicated Wires




Multi-Hop FCoE Solutions

Multi-Hop FCoE is achieved with the support of Virtual E-ports (VE ports) connection two FCFs. Like E_Ports in native FC, VE ports are used to expand the FCoE fabric. VE ports are supported on the Nexus 5000 Series switch as of the NXOS Release 5.0(1)N2(2). There are two options for connecting Nexus 5000 Series switches with the use of VE ports: using single-links or over a port channel. For configuration examples of VE ports, see [Chapter 3, “FCoE Port Configuration Examples.”](#)

In order to maintain fabric isolation, the Cisco Nexus 5000 FCF switches in each fabric should be configured to have the same FC-MAP value. The FC-MAP values should be different between Fabric A and Fabric B. For additional information on FC-MAP configurations, see [Chapter 3, “FCoE Port Configuration Examples.”](#) VE ports brought up between two Cisco Nexus 5000 Series switches with differing FC-MAPs are not supported which ensures that fabrics are not merged by connecting FCFs in Fabric A to FCFs in Fabric B. [Figure 1-10](#) shows FCF connections using VE ports.

Send documentation comments to n5kdocfeedback@cisco.com

Figure 1-10 VE Ports And FCF Mapping



Note VE ports are not supported over vPCs.



Note Port channeling multiple ports in a port channel from an end device to 2232PP FEX and associating it in a VFC is not supported. The end device can be either a Blade Center Chassis (IBM or HP), or a standalone host. This can cause host logins to fail or other unforeseen issues may occur. To workaround this issue, use either a single member port channel bind into a VFC, or a single physical link bind into a VFC. Another possible workaround is to connect the end device directly to the Nexus chassis base ports.

FCoE Operations

This section includes the following topics:

- [Tracking FCoE Statistics, page 1-22](#)
- [SPAN for FC and FCoE Traffic, page 1-23](#)
- [Roles Based Access Control, page 1-25](#)

Tracking FCoE Statistics

FCoE statistics for FCoE traffic transverse an interface on a Cisco Nexus 5000 Series switch can be seen by monitoring the statistics on the vFC interface which is bound to the physical Ethernet interface or port channel interface.

This section includes the following topics:

- [Tracking VE Port Statistics, page 1-23](#)

Send documentation comments to n5kdocfeedback@cisco.com

- [Tracking VF Port Statistics, page 1-23](#)

Tracking VE Port Statistics

The following example shows how to monitor VE port statistics:

```
switch(config-if)# show inter vfc 300
vfc300 is trunking
  Bound interface is port-channel300
  Hardware is Virtual Fibre Channel
  Port WWN is 21:2b:00:05:9b:77:f5:7f
  Admin port mode is E, trunk mode is on
  snmp link state traps are enabled
  Port mode is TE
  Port vsan is 1
  Trunk vsans (admin allowed and active) (3,5)
  Trunk vsans (up) (3,5)
  Trunk vsans (isolated) ()
  Trunk vsans (initializing) ()
  1 minute input rate 15600 bits/sec, 1950 bytes/sec, 21 frames/sec
  1 minute output rate 43664 bits/sec, 5458 bytes/sec, 21 frames/sec
  51295547 frames input, 10484381916 bytes
  0 discards, 0 errors
  39089018 frames output, 10620127132 bytes
  0 discards, 0 errors
  last clearing of "show interface" counters never
  Interface last changed at Mon Jan 17 19:05:27 2011
```

Tracking VF Port Statistics

The following example shows how to monitor VF port statistics:

```
switch(config-if)# show inter vfc 31
vfc31 is trunking (Not all VSANs UP on the trunk)
  Bound interface is Ethernet1/1
  Hardware is Virtual Fibre Channel
  Port WWN is 20:1e:00:05:9b:77:f5:7f
  Admin port mode is F, trunk mode is on
  snmp link state traps are enabled
  Port mode is TF
  Port vsan is 3
  Trunk vsans (admin allowed and active) (3)
  Trunk vsans (up) (3)
  Trunk vsans (isolated) ()
  Trunk vsans (initializing) ()
  1 minute input rate 6912756368 bits/sec, 864094546 bytes/sec, 8640880 frames/sec
  1 minute output rate 6963590568 bits/sec, 870448821 bytes/sec, 396313 frames/sec
  789408333283 frames input, 78940833327276 bytes
  0 discards, 0 errors
  36207053863 frames output, 79510690165704 bytes
  0 discards, 0 errors
  last clearing of "show interface" counters never
  Interface last changed at Mon Jan 17 19:05:21 2011
```

SPAN for FC and FCoE Traffic

This section includes the following topics:

- [Possible SPAN Sources, page 1-24](#)

Send documentation comments to n5kdocfeedback@cisco.com

- [Possible SPAN Destinations, page 1-24](#)
- [SPAN Configuration Examples, page 1-24](#)

Possible SPAN Sources

Following are possible SPAN sources:

- FC interface (only rx-source on 5500 platform)
- VFC interface
- VSAN (not supported on 5500 platform)
- VLAN
- Ethernet interface
- Port channel interface
- SAN port channel interface

Possible SPAN Destinations

Following are possible SPAN destinations:

- FC interface
- Ethernet interface

SPAN Configuration Examples

This example shows how to display configuration information on Ethernet 1/1:

```
switch(config)# show running-config interface eth 1/1
interface Ethernet1/1
    switchport monitor
```

This example shows how to display the health monitoring of all interfaces for failover purposes:

```
switch(config)# show running-config monitor all
monitor session 1 type local
    no description
    source interface vfc33 both
    destination interface Ethernet1/1
    no shut
```

This example shows the health monitoring of session 1:

```
switch(config)# show monitor session 1
    session 1
    -----
type : local
state : up
source intf :
    rx : vfc33
    tx : vfc33
    both : vfc33
source VLANs :
    rx :
source VSANs :
    rx :
destination ports : Eth1/1
```

Send documentation comments to n5kdocfeedback@cisco.com

Legend: f = forwarding enabled, l = learning enabled

This example shows the health monitoring configuration:

```
switch(config)# show running-config monitor
monitor session 1
  source interface fc3/1 tx
  destination interface Ethernet1/1
  no shut
```

This example shows the health monitoring of all sessions:

```
switch(config)# show monitor session all
session 1
-----
type : local
state : up
source intf :
  rx : fc3/1
  tx : fc3/1
  both : fc3/1
source VLANs :
  rx :
source VSANs :
  rx p:
destination ports : Eth1/1
Legend: f = forwarding enabled, l = learning enabled
```

Roles Based Access Control

With the Cisco Nexus Family of switches deploying unified I/O capabilities, the roles of LAN and SAN administrators are converging. To help manage these two different roles on the Cisco Nexus Series Family of switches, the Roles Based Access Control (RBAC) feature facilitates various administrative operations.

When deploying unified I/O within a data center, Cisco recommends defining the following three roles:

- **Unified Administrator**—This role includes all actions that impact both LAN and SAN operations. This role is sometimes referred to as a global administrator.
- **LAN Administrator**—This role includes a set of actions that impact LAN operation while denying any actions that could impact SAN operations.
- **SAN Administrator**—This role includes a set of actions that impact SAN operation while denying any actions that could impact LAN operations.

These are general roles that are used to enforce the operational model where separate LAN and SAN administrative teams retain management control of their perspective networks without interference. More specific roles may be added if operations need to be more tightly defined.

This section includes the following topics:

- [Unified Administrator Role, page 1-26](#)
- [LAN Administrator Role, page 1-26](#)
- [SAN Administrator Role, page 1-26](#)

Send documentation comments to n5kdocfeedback@cisco.com

Unified Administrator Role

The Unified Administrator role may perform all actions. In addition, the Unified Administrator plays a large role in the initial set up of the unified network.

Before implementing a unified network design, the physical interfaces and VLANs used for unified traffic should be identified and defined. Standard implementation of FCoE requires binding a virtual Fibre Channel interface (vFC) to either a physical Ethernet interface or MAC-Address. It is also required to map the VSAN used to carry the FC traffic to a corresponding Ethernet VLAN. While Ethernet interfaces and VLANs normally fall under the scope of a LAN administration, the unified interfaces and FCoE VLANs must be identified so that they can be separated from the LAN administration domain.

Cisco recommends that you identify the interfaces used for Unified I/O, and that you designate a range of VLANs for FCoE use before implementation begins. The Unified Administrator role will configure these unified interfaces and FCoE VLANs.

LAN Administrator Role

This role is assigned all the permissions that impact LAN traffic. This role also denies any actions that would possibly impact SAN traffic (FCoE and FC). One of the main difference between the LAN administrator role and a LAN administrator in a legacy data center without unified I/O is the inability to shut down a physical Ethernet port carrying FCoE traffic. Potentially, both FC and Ethernet traffic could be traveling over the link simultaneously and, therefore, shutting the port could have an impact on SAN operations.

A list of commands which can impact SAN operations, and therefore should be limited from the role of the LAN Administrator, can be found in [Chapter 2, “FCoE and RBAC Configurations.”](#) Individual network designs may require additional limited commands.

SAN Administrator Role

This role is assigned all the permissions that impact SAN traffic. The role also denies actions that would impact LAN traffic.

SAN administration in a unified environment and a legacy SAN environment are similar. Today, unified I/O runs only between the servers and the top-of-rack Cisco Nexus 5000 switch, where FC links are run back into the core of the existing SAN infrastructure. The FC module inside the Cisco Nexus 5000 switch can operate in either NPV or switch mode. The switch most commonly operates in NPV mode and, from a management perspective, looks identical to a FC blade or fabric switch operating in NPV mode.

A list of commands which can impact LAN operations and therefore should be limited from the role of the SAN Administrator can be found in [Chapter 2, “FCoE and RBAC Configurations.”](#) Individual network designs may require additional limited commands.

FCoE Limitations

This section includes the following topics:

- [Generation 1 And Generation 2 CNA Limitations, page 1-27](#)
- [LACP and FCoE To The Host, page 1-27](#)
- [Deploying a Cisco Nexus 5000 Series Switch as an NPIV Core, page 1-27](#)
- [VE Ports on a Cisco Nexus 5010 Switch or Cisco Nexus 5020 Switch, page 1-27](#)

Send documentation comments to n5kdocfeedback@cisco.com

Generation 1 And Generation 2 CNA Limitations

When FCoE was introduced on the Cisco Nexus 5000 Series switch, Cisco worked with QLogic and Emulex to create the first generation of CNA adapters. These CNAs used a pre-standard implementation of the DCBX protocol nicknamed CIN-DCBX. These adapters also did not support the standard FIP implementation as defined in the FCoE Standard (FC-BB-5) and they are often referred to as Pre-FIP adapters.

Starting in 2009, after the ratification of the FCoE standard, second generation CNAs were put out by both QLogic and Emulex that supported standard FIP and FCoE. These CNAs also used a pre-standard version of the DCBX protocol nicknamed CEE-DCBX which has been decided on by multiple vendors to be the de-facto standard until IEEE DCBX is ratified.

Topologies and Platforms Which Require Generation 2 CNAs

While the Cisco Nexus 5010 switch and Nexus 5020 switch are backwards compatible with both Generation 1 and Generation 2 CNAs and support, the Nexus 2000 Fabric Extenders and the Nexus 5500 Platform switches only support Generation 2 CNA connections. Also, Generation 2 CNAs are required when connecting a host using vPC into a fabric, whether the host is running FCoE or just native Ethernet.

LACP and FCoE To The Host

Today, when deploying FCoE over a host-facing vPC, the vFC interface is bound to the port channel interfaces associated with the vPC. This requires that the port channel interface be up and forwarding before FCoE traffic can be switched. Cisco recommends when running vPC in an Ethernet environment is to use LACP in order to negotiate the parameters on both sides of the port channel to ensure that configurations between both sides is consistent.

However, if there are inconsistencies in any of the Ethernet configuration parameters LACP uses to bring up the port channel interface, both sides of the virtual port channel will remain down. This means that FCoE traffic from the host is now dependent on the correct configuration on the LAN/Ethernet side. When this dependency occurs, Cisco recommends that you use the static port channel configuration (channel-group # mode on) when deploying vPC and FCoE to the same host.

Deploying a Cisco Nexus 5000 Series Switch as an NPIV Core

The Nexus 5000 Series switch supports both NPV and NPIV functionality. If acting as an NPIV core switch with downstream NPV switches attached to it, it is important to note that hosts and targets which are communicating to one another can not be attached to the same downstream NPV device.

VE Ports on a Cisco Nexus 5010 Switch or Cisco Nexus 5020 Switch

Cisco Nexus 5000 Series and Cisco Nexus 5500 Platform switches support VE port connections. On Cisco Nexus 5010 and Nexus 5020 switches, VE ports can be configured between two switches using a single port channel or multiple individual links. VE ports configured between two switches using multiple port channels is not supported. This has to do with the number of MAC addresses available for the VE port on the Cisco Nexus 5010 switch and the Cisco Nexus 5020 switch. This limitation does not apply to the Cisco Nexus 5500 Platform.

Send documentation comments to n5kdocfeedback@cisco.com

Additional Information

See “Configuring FCoE NPV” in the *Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide*:

http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html

Cisco Nexus 5000 Series Switch overview information:

<http://www.cisco.com/en/US/products/ps9670/index.html>

Cisco Nexus 5000 Series Configuration Guides:

http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html

Fibre Channel over Ethernet information: www.fcoe.com



CHAPTER 2

FCoE and RBAC Configurations

This chapter includes information about RBAC configurations in relation to FCoE operations and it includes the following sections:

- [Global Administrator Actions, page 2-1](#)
- [LAN Administrator Actions, page 2-1](#)
- [SAN Administrator Actions, page 2-4](#)
- [Sample Configurations, page 2-6](#)

Global Administrator Actions

The Global administrator role is unrestricted and all commands are available.

LAN Administrator Actions

This section lists the commands that the LAN administrator may not perform. Commands that are not listed are implicitly permitted.

Global Level Deny Actions

```
switch(config)# feature lACP
switch(config)# feature tacacs+
switch(config)# feature udld
switch(config)# feature fcoe
switch(config)# aaa *
switch(config)# boot *
switch(config)# cfs *
switch(config)# class-map *
switch(config)# device-alias *
switch(config)# diagnostic *
switch(config)# fex *
switch(config)# hw-module logging onboard *
switch(config)# license *
switch(config)# line *
switch(config)# lldp *
switch(config)# monitor session *
switch(config)# ntp *
switch(config)# policy-map *
switch(config)# privilege *
switch(config)# radius-server *
```

Send documentation comments to n5kdocfeedback@cisco.com

```
switch(config)# role *
switch(config)# snmp-server *
switch(config)# spanning-tree
    bridge assurance *
    loopguard *
    mode *
    mst *
    pathcost *
    port type *
    vlan <fcoe-vlan>
switch(config)# ssh *
switch(config)# system
    core *
    default switchport *
    jumbomtu *
    qos *
switch(config)# tacacs+ *
switch(config)# telnet server enable
switch(config)# trunk protocol enable
switch(config)# username *
switch(config)# vrf *
switch(config)# xml server *
```

This section includes the following topics:

- [VLAN-Level Deny Actions, page 2-2](#)
- [Interface-Level Deny Actions, page 2-2](#)
- [FC Deny Actions, page 2-3](#)

VLAN-Level Deny Actions

Deny Actions for All VLANs

```
switch(config)# vlan vlan
switch(config-vlan)# fcoe
```

Deny Actions for Pre-determined FCoE VLANs

```
switch(config)# no vlan fcoe-vlan
switch(config)# vlan fcoe-vlan *
switch(config)# spanning-tree vlan fcoe-vlan *
switch(config-mst)# instance n vlan fcoe-vlan
switch(config)# mac-address-table aging-time t vlan fcoe-vlan
switch(config)# mac-address-table static aaaa.bbbb.cccc vlan fcoe-vlan
switch(config-monitor)# source vlan fcoe-vlan
switch(config)# vlan fcoe-vlan
switch(config-vlan)# ip igmp snooping *
```

Interface-Level Deny Actions

Interface-Level Deny Actions



Note

Access to the management interface is limited to the unified administrator.

```
switch(config)# interface mgmt *
```


Send documentation comments to n5kdocfeedback@cisco.com

Deny Actions for Pre-determined Ethernet Interfaces Designated to Carry FCoE Traffic

```
switch(config-if)# bandwidth *
switch(config-if)# fcoe *
switch(config-if)# flowcontrol *
switch(config-if)# link debounce *
switch(config-if)# lldp *
switch(config-if)# priority-flow-control *
switch(config-if)# service-policy *
switch(config-if)# shutdown
switch(config-if)# shutdown force
switch(config-if)# spanning-tree bpdudfilter
switch(config-if)# spanning-tree bpduguard
switch(config-if)# spanning-tree cost *
switch(config-if)# spanning-tree guard *
switch(config-if)# spanning-tree link-type *
switch(config-if)# spanning-tree mst *
switch(config-if)# spanning-tree port type *
switch(config-if)# spanning-tree port-priority *
switch(config-if)# speed *
switch(config-if)# switchport host
switch(config-if)# switchport mode *
switch(config-if)# switchport monitor
switch(config-if)# switchport trunk native vlan <fcoe-vlan>
switch(config-if)# switchport trunk allowed vlan <range>
switch(config-if)# switchport trunk allowed vlan add <fcoe-vlan>
switch(config-if)# switchport trunk allowed vlan all
switch(config-if)# switchport trunk allowed vlan except *
switch(config-if)# switchport trunk allowed vlan none
switch(config-if)# switchport trunk allowed vlan remove <fcoe-vlan>
```

FC Deny Actions

FC Deny Actions



Note

The LAN administrator may not execute SAN-related commands.

```
switch(config)# fabric-binding *
switch(config)# fcalias *
switch(config)# fcdomain *
switch(config)# fcdroplacency *
switch(config)# fcflow *
switch(config)# fcid-allocation *
switch(config)# fcinterop *
switch(config)# fcns *
switch(config)# fcroute *
switch(config)# fcs *
switch(config)# fcsp *
switch(config)# fctimer *
switch(config)# fdmi *
switch(config)# fspf *
switch(config)# in-order-guarantee
switch(config)# interface fc *
switch(config)# interface san-port-channel *
switch(config)# interface vfc *
switch(config)# npiv *
switch(config)# npv *
switch(config)# port-security enable
switch(config)# port-track enable
switch(config)# rib *
```

Send documentation comments to n5kdocfeedback@cisco.com

```
switch(config)# rlr *
switch(config)# rscn *
switch(config)# scsi-target *
switch(config)# system default zone *
switch(config)# vsan database *
switch(config)# wwn *
switch(config)# zone *
switch(config)# zoneset *
```

SAN Administrator Actions

This section lists the commands that the SAN administrator may not perform. Commands that are not listed are implicitly permitted.

Global Level Deny Actions

```
switch(config)# feature * (except feature fcoe)
switch(config)# aaa *
switch(config)# boot *
switch(config)# cfs *
switch(config)# class-map *
switch(config)# device-alias *
switch(config)# diagnostic *
switch(config)# fex *
switch(config)# hw-module logging onboard *
switch(config)# ip *
switch(config)# ipv6 *
switch(config)# license *
switch(config)# line *
switch(config)# lldp *
switch(config)# mac-address-table *
switch(config)# monitor session *
switch(config)# ntp *
switch(config)# policy-map *
switch(config)# privilege *
switch(config)# radius-server *
switch(config)# role *
switch(config)# snmp-server *
switch(config)# spanning-tree
    bridge assurance *
    loopguard *
    mode *
    mst *
    pathcost *
    port type *
    vlan <non-fcoe-vlan>
switch(config)# ssh *
switch(config)# system
    core *
    default switchport *
    jumbomt *
    qos *
switch(config)# tacacs+ *
switch(config)# telnet server enable
switch(config)# trunk protocol enable
switch(config)# username *
switch(config)# vrf *
switch(config)# xml server *
```

This section includes the following topics:

Send documentation comments to n5kdocfeedback@cisco.com

- [VLAN Level Deny Actions, page 2-5](#)
- [Interface Level Deny Actions, page 2-5](#)
- [LAN Deny Actions, page 2-6](#)

VLAN Level Deny Actions

Deny Actions for Pre-determined Non-FCoE VLANs

```
switch(config)# no vlan <non-fcoe-vlan>
switch(config)# vlan <non-fcoe-vlan>*
switch(config)# spanning-tree vlan <non-fcoe-vlan>*
switch(config-mst)# instance n vlan <non-fcoe-vlan>
switch(config)# mac-address-table aging-time t vlan <non-fcoe-vlan>
switch(config)# mac-address-table static aaaa.bbbb.cccc vlan <non-fcoe-vlan>
switch(config-monitor)# source vlan <non-fcoe-vlan>
switch(config)# vlan <non-fcoe-vlan>
switch(config-vlan)# ip igmp snooping *
switch(config-if)# spanning-tree vlan <non-fcoe-vlan>
```

Interface Level Deny Actions

Interface Level Deny Actions



Note

Access to the management interface is limited to the unified administrator.

```
switch (config)# interface mgmt *
```

Deny Actions for Pre-determined Ethernet Interfaces Designated to not Carry FCoE Traffic

The SAN administrator may execute **no** commands on these interfaces.

Deny Actions for Pre-determined Ethernet Interfaces Designated to Carry FCoE Traffic

This deny-list applies to Ethernet, port-channel, and vEthernet interfaces that are designated to carry FCoE traffic.

```
switch(config-if)# bandwidth *
switch(config-if)# fcoe *
switch(config-if)# flowcontrol *
switch(config-if)# link debounce *
switch(config-if)# lldp *
switch(config-if)# priority-flow-control *
switch(config-if)# service-policy *
switch(config-if)# shutdown
switch(config-if)# shutdown force
switch(config-if)# shutdown lan // TBD. This is a new command to shut stop LAN VLANs
switch(config-if)# spanning-tree bpduguard
switch(config-if)# spanning-tree bpduguard
switch(config-if)# spanning-tree cost *
switch(config-if)# spanning-tree guard *
switch(config-if)# spanning-tree link-type *
switch(config-if)# spanning-tree mst *
switch(config-if)# spanning-tree port type *
switch(config-if)# spanning-tree port-priority *
switch(config-if)# speed *
switch(config-if)# switchport host
```

Send documentation comments to n5kdocfeedback@cisco.com

```
switch(config-if)# switchport mode *
switch(config-if)# switchport monitor
switch(config-if)# switchport trunk native *
switch(config-if)# switchport trunk allowed vlan <range>
switch(config-if)# switchport trunk allowed vlan add <non-fcoe-vlan>
switch(config-if)# switchport trunk allowed vlan all
switch(config-if)# switchport trunk allowed vlan except *
switch(config-if)# switchport trunk allowed vlan none
switch(config-if)# switchport trunk allowed vlan remove <non-fcoe-vlan>
```

LAN Deny Actions

LAN Deny Actions

The SAN administrator can not execute LAN-related commands.

```
switch(config)# cdp *
switch(config)# ip igmp snooping *
switch(config)# port-channel load-balance ethernet
switch(config)# rmon
switch(config)# track
```

Sample Configurations

The following configurations are used to create both LAN and SAN administrative roles. These configurations follow the outline listed above concerning the commands that are assigned to or withheld from each role. Configuration is not needed for the Global Administrator who automatically has access to all configuration commands.



Note

This configuration assumes that vFC 1 is mapped to Ethernet 1/1 and that VLAN 100 has been designated the FCoE VLAN. This configuration is based on the specific environment and which Ethernet ports and VLANs have been pre-determined to carry FCoE traffic.

LAN-Admin Configuration

```
role name LAN-admin
description assume vlan 100 is fcoe enabled and eth1/1 is an vfc bound (fcoe) interface
rule 97 deny command config t ; feature lacp
rule 96 deny command config t ; feature tacacs+
rule 95 deny command config t ; feature uddl
rule 94 deny command config t ; feature fcoe
rule 93 deny command config t ; aaa *
rule 92 deny command config t ; boot *
rule 91 deny command config t ; cfs *
rule 90 deny command config t ; class-map *
rule 89 deny command config t ; device-alias *
rule 88 deny command config t ; diagnostic *
rule 87 deny command config t ; fex *
```

Send documentation comments to n5kdocfeedback@cisco.com

```
rule 86 deny command config t ; hw-module logging onboard *
rule 85 deny command config t ; license *
rule 84 deny command config t ; line *
rule 83 deny command config t ; lldp *
rule 82 deny command config t ; monitor session *
rule 81 deny command config t ; ntp *
rule 80 deny command config t ; policy-map *
rule 79 deny command config t ; privilege *
rule 78 deny command config t ; radius-server *
rule 77 deny command config t ; role *
rule 76 deny command config t ; snmp-server *
rule 75 deny command config t ; ssh *
rule 74 deny command config t ; system *
rule 73 deny command config t ; no system *
rule 72 deny command config t ; tacacs+ *
rule 71 deny command config t ; telnet server enable
rule 70 deny command config t ; trunk protocol enable
rule 69 deny command config t ; username *
rule 68 deny command config t ; vrf *
rule 67 deny command config t ; xml server *
rule 66 deny command config t ; fabric-binding *
rule 65 deny command config t ; fcalias *
rule 64 deny command config t ; fcdomain *
rule 63 deny command config t ; fcdroplatency *
rule 62 deny command config t ; fcflow *
rule 61 deny command config t ; fcid-allocation *
rule 60 deny command config t ; fcinterop *
rule 59 deny command config t ; fcns *
rule 58 deny command config t ; fcroute *
rule 57 deny command config t ; fcs *
rule 56 deny command config t ; fcsp *
rule 55 deny command config t ; fctimer *
rule 54 deny command config t ; fdmi *
rule 53 deny command config t ; fspf *
rule 52 deny command config t ; in-order-guarantee
rule 51 deny command config t ; npiv *
rule 50 deny command config t ; npv *
rule 49 deny command config t ; port-security enable
```

Send documentation comments to n5kdocfeedback@cisco.com

```

rule 48 deny command config t ; port-track enable
rule 47 deny command config t ; rib *
rule 46 deny command config t ; rlir *
rule 45 deny command config t ; rscn *
rule 44 deny command config t ; scsi-target *
rule 43 deny command config t ; vsan database *
rule 42 deny command config t ; wwn *
rule 41 deny command config t ; zone *
rule 40 deny command config t ; zoneset *
rule 39 deny command config t ; vlan * ; fcoe *
rule 38 deny command config t ; vlan * ; no fcoe *
rule 37 deny command config t ; spanning-tree vlan 100
rule 36 permit command config t ; spanning-tree vlan *
rule 35 deny command config t ; spanning-tree *
rule 34 deny command config t ; mac-address-table aging-time * vlan 100
rule 33 deny command config t ; mac-address-table static * vlan 100 *
rule 32 deny command config t ; monitor session * ; source vlan 100
rule 31 deny command config t ; vlan 100 *
rule 30 deny command config t ; no vlan 100 *
rule 29 deny command config t ; interface Ethernet1/1 ; bandwidth *
rule 28 deny command config t ; interface Ethernet1/1 ; fcoe *
rule 27 deny command config t ; interface Ethernet1/1 ; flowcontrol *
rule 26 deny command config t ; interface Ethernet1/1 ; link debounce *
rule 25 deny command config t ; interface Ethernet1/1 ; lldp *
rule 24 deny command config t ; interface Ethernet1/1 ; priority-flow-control *
rule 23 deny command config t ; interface Ethernet1/1 ; service-policy *
rule 22 deny command config t ; interface Ethernet1/1 ; shutdown
rule 21 deny command config t ; interface Ethernet1/1 ; shutdown force
rule 20 deny command config t ; interface Ethernet1/1 ; spanning-tree bpdudfilter *
rule 19 deny command config t ; interface Ethernet1/1 ; spanning-tree bpduguard *
rule 18 deny command config t ; interface Ethernet1/1 ; spanning-tree cost *
rule 17 deny command config t ; interface Ethernet1/1 ; spanning-tree guard *
rule 16 deny command config t ; interface Ethernet1/1 ; spanning-tree link-type *
rule 15 deny command config t ; interface Ethernet1/1 ; spanning-tree mst *
rule 14 deny command config t ; interface Ethernet1/1 ; spanning-tree port type *
rule 13 deny command config t ; interface Ethernet1/1 ; spanning-tree port-priority *
rule 12 deny command config t ; interface Ethernet1/1 ; speed *
rule 11 deny command config t ; interface Ethernet1/1 ; switchport host

```

Send documentation comments to n5kdocfeedback@cisco.com

```
rule 10 deny command config t ; interface Ethernet1/1 ; switchport mode *
rule 9 deny command config t ; interface Ethernet1/1 ; switchport monitor
rule 8 deny command config t ; interface Ethernet1/1 ; switchport trunk native vlan 100
rule 7 deny command config t ; interface Ethernet1/1 ; switchport trunk allowed vlan *
rule 6 deny command config t ; interface Ethernet1/1 ; switchport trunk allowed vlan add 100
rule 5 deny command config t ; interface Ethernet1/1 ; switchport trunk allowed vlan all
rule 4 deny command config t ; interface Ethernet1/1 ; switchport trunk allowed vlan except *
rule 3 deny command config t ; interface Ethernet1/1 ; switchport trunk allowed vlan none
rule 2 deny command config t ; interface Ethernet1/1 ; switchport trunk allowed vlan remove 100
rule 1 permit read-write
interface policy deny
    permit interface eth1/1-40
vlan policy deny
    permit vlan 100-200
vsan policy deny
```

SAN-Admin Configuration

```
role name SAN-admin
description assuming vlan 100 is fcoe enabled and vfc1 has been bound to eth1/1
rule 83 permit command config t ; vlan * ; fcoe *
rule 82 deny command config t ; vlan * ; *
rule 81 deny command config t ; ip igmp snooping *
rule 80 deny command config t ; cdp *
rule 79 deny command config t ; port-channel load-balance ethernet *
rule 78 deny command config t ; rmon *
rule 77 deny command config t ; track *
rule 76 deny command config t ; no ip igmp *
rule 75 deny command config t ; no cdp *
rule 74 deny command config t ; no port-channel load-balance *
rule 73 deny command config t ; no rmon *
rule 72 deny command config t ; no track *
rule 71 deny command config t ; interface * ; switchport trunk native *
rule 70 deny command config t ; interface * ; switchport trunk allowed vlan *
rule 69 deny command config t ; interface * ; switchport trunk allowed vlan add 100
rule 68 deny command config t ; interface * ; switchport trunk allowed vlan all
rule 67 deny command config t ; interface * ; switchport trunk allowed vlan except *
rule 66 deny command config t ; interface * ; switchport trunk allowed vlan none
rule 65 deny command config t ; interface * ; switchport trunk allowed vlan remove 100
```

Send documentation comments to n5kdocfeedback@cisco.com

```
rule 64 deny command config t ; interface * ; bandwidth *
rule 63 deny command config t ; interface * ; fcoe *
rule 62 deny command config t ; interface * ; flowcontrol *
rule 61 deny command config t ; interface * ; link debounce *
rule 60 deny command config t ; interface * ; lldp *
rule 59 deny command config t ; interface * ; priority-flow-control *
rule 58 deny command config t ; interface * ; service-policy *
rule 57 deny command config t ; interface * ; shutdown
rule 56 deny command config t ; interface * ; shutdown force
rule 55 deny command config t ; interface * ; shutdown lan
rule 54 deny command config t ; interface * ; spanning-tree bpdudfilter
rule 53 deny command config t ; interface * ; spanning-tree bpduguard
rule 52 deny command config t ; interface * ; spanning-tree cost *
rule 51 deny command config t ; interface * ; spanning-tree guard *
rule 50 deny command config t ; interface * ; spanning-tree link-type *
rule 49 deny command config t ; interface * ; spanning-tree mst *
rule 48 deny command config t ; interface * ; spanning-tree port type *
rule 47 deny command config t ; interface * ; spanning-tree port-priority *
rule 46 deny command config t ; interface * ; speed *
rule 45 deny command config t ; interface * ; switchport host
rule 44 deny command config t ; interface * ; switchport mode *
rule 43 deny command config t ; interface * ; switchport monitor
rule 42 deny command config t ; no vlan 100 *
rule 41 permit command config t ; feature fcoe
rule 40 deny command config t ; feature *
rule 39 deny command config t ; aaa *
rule 38 deny command config t ; boot *
rule 37 deny command config t ; cfs *
rule 36 deny command config t ; class-map *
rule 35 deny command config t ; device-alias *
rule 34 deny command config t ; diagnostic *
rule 33 deny command config t ; fex *
rule 32 deny command config t ; hw-module logging onboard *
rule 31 deny command config t ; ip *
rule 30 deny command config t ; ipv6 *
rule 29 deny command config t ; license *
rule 28 deny command config t ; line *
rule 27 deny command config t ; lldp *
```


Send documentation comments to n5kdocfeedback@cisco.com

```
rule 26 deny command config t ; mac-address-table *
rule 25 deny command config t ; monitor session *
rule 24 deny command config t ; ntp *
rule 23 deny command config t ; policy-map *
rule 22 deny command config t ; privilege *
rule 21 deny command config t ; radius-server *
rule 20 deny command config t ; role *
rule 19 deny command config t ; snmp-server *
rule 18 deny command config t ; spanning-tree bridge assurance *
rule 17 deny command config t ; spanning-tree loopguard *
rule 16 deny command config t ; spanning-tree mode *
rule 15 deny command config t ; spanning-tree mst *
rule 14 deny command config t ; spanning-tree pathcost *
rule 13 deny command config t ; spanning-tree port type *
rule 12 deny command config t ; ssh *
rule 11 deny command config t ; system core *
rule 10 deny command config t ; system default switchport *
rule 9 deny command config t ; system jumbomtu *
rule 8 deny command config t ; system qos *
rule 7 deny command config t ; tacacs+ *
rule 6 deny command config t ; telnet server enable
rule 5 deny command config t ; trunk protocol enable
rule 4 deny command config t ; username *
rule 3 deny command config t ; vrf *
rule 2 deny command config t ; xml server *
rule 1 permit read-write
vlan policy deny
  permit vlan 100-100
interface policy deny
  permit interface fc3/1-4
  permit interface Ethernet1/1
  permit interface vfc1
```

Send documentation comments to n5kdocfeedback@cisco.com



CHAPTER 3

FCoE Port Configuration Examples

This appendix describes port configuration examples relating to FCoE topologies and it includes the following sections:

- [VE Port Configuration Example, page 3-1](#)
- [FCoE VE Port Topology Example, page 3-1](#)
- [Enabling FCoE and Verifying QoS Configuration, page 3-2](#)
- [Configuring VE Ports, page 3-5](#)

VE Port Configuration Example

This section provides a sample configuration of the Cisco Nexus 5000 Series switch FCoE VE Port implementation. The configuration covers the switches in switch mode. FCoE initiators are used in this lab. You can attach either FC F Port storage directly to a Nexus 5000 Series switch FC GEMs, or use an FCoE target.



Note

This example can be used for configuring VE ports between two Cisco Nexus 5000 Series switches in both fabrics. It does not include server configurations.

FCoE VE Port Topology Example


[Figure 3-1](#) shows the topology that was used for the configuration example. The following configuration parameters are used in this topology:

- FCoE VLAN for Fabric A: 10
- FCoE VSAN for Fabric A: 10
- FCoE VLAN for Fabric B: 20
- FCoE VSAN for Fabric B: 20
- Ethernet Only VLAN across both fabrics: 200

You should choose these values before the time of configuration.

Send documentation comments to n5kdocfeedback@cisco.com

Figure 3-1 FCoE VE Port Topology



Note

The FCoE VLAN/VSAN numbering does not have to be the same within the fabric. As a best practice, use different FCoE VLANs and VSAN numbers between the two fabrics to avoid confusion. Configurations have often been set up to assign ODD VLAN/VSANs for one fabric and EVEN VLANs/VSANs for the other fabric. This is just one example of keeping the numbers separate between the two fabrics

Enabling FCoE and Verifying QoS Configuration

Step 1 Enable FCoE.

```
switch# configure terminal
switch(config)# feature fcoe
FC license checked out successfully
fc_plugin extracted successfully
FC plugin loaded successfully
FCoE manager enabled successfully
FC enabled on all modules successfully
```

Step 2 (Optional) If you do not want to use the default Quality of Service (QoS) settings, specify your own policies:

Send documentation comments to n5kdocfeedback@cisco.com



Note Note: if you use custom policies, **class-fcoe** must be included in your QoS policies.

```
switch(config) system qos
switch(config-sys-qos)# service-policy type qos input fcoe-customized-in-policy-name
switch(config-sys-qos)# service-policy type queuing input
fcoe-customized-in-policy-name
switch(config-sys-qos)# service-policy type queuing output
fcoe-customized-out-policy-name
switch(config-sys-qos)# service-policy type network-qos fcoe-customized-nq-policy-name
```

Step 3 Verify that the FCoE policy maps can be found in the running configuration:



Note Note: If you specified customized QoS policy map names in [Step 2](#), make sure you replace the default map names with your customized map names.

```
switch(config-sys-qos)# show policy-map system

Type network-qos policy-maps
=====

policy-map type network-qos system
  class type network-qos class-fcoe
    match qos-group 1

    pause no-drop
    mtu 2158
  class type network-qos class-default
    match qos-group 0

    mtu 1500

Service-policy (qos) input:    system
policy statistics status:    disabled

Class-map (qos):    class-fcoe (match-any)
  Match: cos 3
  set qos-group 1

Class-map (qos):    class-default (match-any)
  Match: any
  set qos-group 0

Service-policy (queuing) input:  default-in-policy
policy statistics status:  disabled

Class-map (queuing):  class-fcoe (match-any)
  Match: qos-group 1
  bandwidth percent 50

Class-map (queuing):  class-default (match-any)
  Match: qos-group 0
  bandwidth percent 50

Service-policy (queuing) output:  default-out-policy
policy statistics status:  disabled

Class-map (queuing):  class-fcoe (match-any)
  Match: qos-group 1
  bandwidth percent 50
```

Send documentation comments to n5kdocfeedback@cisco.com

```
Class-map (queuing): class-default (match-any)
  Match: qos-group 0
  bandwidth percent 50
```

Quality of Service configuration on the Nexus 5000 series consists of three main constructs:

- Class-map and policy-map type qos: for classification purposes
- Class-map and policy-map type network: for network properties such as drop and no drop, queue size
- Class-map and policy-map type queueing: for bandwidth allocation

This exercise consists of changing the bandwidth allocation and the COS settings for FCoE.

Without proper configuration of class-fcoe in QoS, the following problems may occur:

- vFC interfaces do not come up (CNAs require advertisement of DCB parameters for FCoE)
- Drops noticed for I/Os



Note

QoS has the following guidelines:

- A classification policy-map only applies in input
- A network policy-map applies globally (system)
- A queueing policy-map normally is meaningful in output, but since the exercise uses it to control the bandwidth allocation from CNA to the Cisco Nexus 5000 Series switch, in this case it is applied in input

Beginning in Cisco NX-OS Release 5.0(2)N1(1), you can modify the buffer allocation for no-drop classes:

```
switch(config-pmap-nq)# policy-map type network-qos nqos_policy
switch(config-pmap-nq)# class type network-qos nqos_class
switch(config-pmap-nq-c)# pause no-drop buffer-size <size> pause-threshold <threshold>
resume-threshold <threshold>
```

Step 4 Verify the FCoE system class is active:

```
switch(config-sys-qos)# show queuing interface ethernet 1/1
Ethernet1/1 queuing information:
TX Queuing
  qos-group sched-type oper-bandwidth
  0 WRR 50
  1 WRR 50
RX Queuing
  qos-group 0
  q-size: 370240, HW MTU: 1500 (1500 configured)
  drop-type: drop, xon: 0, xoff: 2314
  Statistics:
  Pkts received over the port : 0
  Ucastpkts sent to the cross-bar : 0
  Mcastpkts sent to the cross-bar : 0
  Ucastpkts received from the cross-bar : 0
  Pkts sent to the port : 0
  Pkts discarded on ingress : 0
  Per-priority-pause status : Rx (Inactive), Tx (Inactive)
  qos-group 1
  q-size: 79360, HW MTU: 2158 (2158 configured)
  drop-type: no-drop, xon: 128, xoff: 252
```

Send documentation comments to n5kdocfeedback@cisco.com

```

Statistics:
Pkts received over the port : 0
Ucastpkts sent to the cross-bar : 0
Mcastpkts sent to the cross-bar : 0
Ucastpkts received from the cross-bar : 0
Pkts sent to the port : 0
Pkts discarded on ingress : 0
Per-priority-pause status : Rx (Inactive), Tx (Inactive)
Total Multicast crossbar statistics:
Mcastpkts received from the cross-bar : 0

```

- Step 5** Repeat [Step 1](#) through [Step 4](#) on both upstream Cisco Nexus 5000 Series switches (CORE_N5k-1 and CORE_N5k-2 in this example).

Configuring VE Ports

FCoE VLAN and VSAN numbering in this example is as follows:

- Fabric A uses FCoE VLAN 10 and VSAN 10
- Fabric B uses FCoE VLAN 20 and VSAN 20



Note

There are two switches in Fabric A and two switches in Fabric B. The FCoE VLAN/VSANs must match between the switches in the same fabric in order to bring up the VE port between them.

- Step 1** Configure the VSAN on the Nexus 5000 Series switch for Fabric A:

```

switch(config)#
switch(config)# vsan database
switch(config-vsan-db)# vsan 10

```

- Step 2** Configure the FCoE VLAN to VSAN mapping and verify that it is up and operational for Fabric A:

```

switch(config)# vlan 10
switch(config-vlan)# fcoe vsan 10
switch(config-vlan)#
switch(config-vlan)# show vlan fcoe
Original VLAN ID      Translated VSAN ID      Association State
-----
10                    10 Operational
switch(config-vlan)#

```

- Step 3** Repeat [Step 1](#) and [Step 2](#) on the upstream Nexus 5000 Series switch in Fabric A.

- Step 4** Configure the VSAN on the Nexus 5000 for Fabric B:

```

switch(config)#
switch(config)# vsan database
switch(config-vsan-db)# vsan 20

```

- Step 5** Configure the FCoE VLAN to VSAN mapping and verify that it is up and operational for Fabric B

```

switch(config)# vlan 20
switch(config-vlan)# fcoe vsan 20
switch(config-vlan)#
switch(config-vlan)# show vlan fcoe
Original VLAN ID      Translated VSAN ID      Association State
-----
20                    20 Operational
switch(config-vlan)#

```

Send documentation comments to n5kdocfeedback@cisco.com

- Step 6** Repeated [Step 1](#) and [Step 2](#) on the upstream Nexus 5000 Series switch in Fabric B.
- Step 7** Configure the underlying 10-Gigabit Ethernet port that the vFC interface will be bound to. The VE port will use this interface as the physical transport for FCoE traffic between the two switches. This interface needs to be configured to trunk the appropriate FCoE VLAN as well as the Ethernet VLAN (in this example, we are using VLAN 200 to carry Ethernet traffic).

The 10-Gigabit Ethernet interfaces connecting the switches in this lab are shown in the topology above:

- Fabric A uses FCoE VLAN 10 and VSAN 10
- Fabric B uses FCoE VLAN 20 and VSAN 20
- PODX-N5K-1 (Fabric A) uses Ethernet 1/15 to connect to CORE N5K1
- PODX-N5K-2 (Fabric B) uses Ethernet 1/16 to connect to CORE N5K2

Configuration for both switches in Fabric A:

```
switch(config)# vlan 200
switch(config)# interface ethernet 1/15
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 10, 200
switch(config-if)#
```

Configuration for both switches in Fabric B:

```
switch(config)# vlan 200
switch(config)# interface ethernet 1/16
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 20, 200
switch(config-if)#
```

- Step 8** Configure the vFC interface on the switch that will be bound to the VE port and add this vFC interface to VSAN 44 in the VSAN database:

The vFC numbers for the VE ports are as follows:

- Fabric A uses FCoE VLAN 10 and VSAN 10
- Fabric B uses FCoE VLAN 20 and VSAN 20
- POD1-N5K-1 (Fabric A) uses Ethernet 1/15 to connect to CORE N5K1
- POD1-N5K-2 (Fabric B) uses Ethernet 1/16 to connect to CORE N5K2
- POD1-N5K-1 (Fabric A) uses vfc 15 and binds it to Ethernet 1/15
- POD1-N5K-2 (Fabric B) uses vfc 16 and binds it to Ethernet 1/16

Configuration for both switches in Fabric A:

```
switch(config)# int vfc 15
switch(config-if)# switchport mode e
switch(config-if)# switchport trunk allowed vsan 10
switch(config-if)# bind interface eth 1/15
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)# vsan database
switch(config-vsan-db)# vsan 10 interface vfc 15
switch(config-vsan-db)# show vsan membership
vsan 1 interfaces:
fc2/1          fc2/2          fc2/3          fc2/4
fc2/5          fc2/6          fc2/7          fc2/8
vsan 10 interfaces:
vfc15
vsan 4079(evfp_isolated_vsan) interfaces:
```


Send documentation comments to n5kdocfeedback@cisco.com

```
vsan 4094(isolated_vsan) interfaces:
switch(config-vsan-db)# exit
```

Configuration for both switches in Fabric B:

```
switch(config)# int vfc 16
switch(config-if)# switchport mode e
switch(config-if)# switchport trunk allowed vsan 20
switch(config-if)# bind interface eth 1/16
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)# vsan database
switch(config-vsan-db)# vsan 20 interface vfc 16
switch(config-vsan-db)# show vsan membership
vsan 1 interfaces:
fc2/1          fc2/2          fc2/3          fc2/4
fc2/5          fc2/6          fc2/7          fc2/8
vsan 20 interfaces:
vfc16
vsan 4079(evfp_isolated_vsan) interfaces:
vsan 4094(isolated_vsan) interfaces:
switch(config-vsan-db)# exit
```



Note

Don't forget that these interface configurations must be configured on both sides of the ISL connecting the two switches in the same fabric.

- Step 9** Verify that the vFC is up and operational. By default, the vFC will show as trunking. Make sure that it is bound to the correct physical interface and that VSAN 44 shows as allowed and active as well as up on the vFC interface.

Verify both switches in Fabric A:

```
switch(config)# show int vfc 15
vfc15 is trunking
Bound interface is Ethernet1/15
Hardware is Virtual Fibre Channel
Port WWN is 20:0e:00:0d:ec:b4:43:7f
Peer port WWN is 00:00:00:00:00:00:00:00
Admin port mode is E, trunk mode is on
snmp link state traps are enabled
Port mode is TE
Port vsan is 10
Trunk vsans (admin allowed and active) (10)
Trunk vsans (up) (10)
Trunk vsans (isolated) ()
Trunk vsans (initializing) ()
1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
13 frames input, 1028 bytes
0 discards, 0 errors
13 frames output, 1180 bytes
0 discards, 0 errors
last clearing of "show interface" counters never
Interface last changed at Sat Nov 6 17:58:39 2010
```

Verify both switches in Fabric A:

```
switch(config)# show int vfc 16
vfc16 is trunking
Bound interface is Ethernet1/16
Hardware is Virtual Fibre Channel
Port WWN is 20:0e:00:0d:ec:b4:43:7d
Peer port WWN is 00:00:00:00:00:00:00:00
```

Send documentation comments to n5kdocfeedback@cisco.com

```
Admin port mode is E, trunk mode is on
snmp link state traps are enabled
Port mode is TE
Port vsan is 20
Trunk vsans (admin allowed and active) (20)
Trunk vsans (up) (20)
Trunk vsans (isolated) ()
Trunk vsans (initializing) ()
1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 13 frames input, 1028 bytes
 0 discards, 0 errors
 13 frames output, 1180 bytes
 0 discards, 0 errors
last clearing of "show interface" counters never
Interface last changed at Sat Nov  6 17:58:39 2010
```



CHAPTER 4


FCoE with vPC Configuration Example

Beginning with Cisco NX-OS Release 4.1(3)N1(1), the Cisco Nexus 5000 Series switch supports vPCs which can be configured to increase bandwidth and increased load-balancing to the Ethernet fabric. This appendix includes a sample configuration on how to configure FCoE when using vPCs on the Cisco Nexus 5000 Series switch and includes the following sections:

- [Cisco Nexus 5000 Series Switch vPC Configuration Example, page 4-2](#)
- [Cisco Nexus 5000 Series Switch FCoE Configuration Example, page 4-5](#)

Figure 4-1 shows the topology used in the examples described in this appendix.

Figure 4-1 Nexus 5000 FCoE and vPC Lab Topology



Note

FCoE VLANs should not be trunked across vPC peer-links.

The configuration example includes the following parameters:

```
switchname: n5k-tme-1
switchname: n5k-tme-2
mgmt ip: 172.25.182.66
mgmt ip: 172.25.182.67
```

The configuration example includes the following hardware:

- Dell Server PE2950
- QLogic QLE8142 (Schultz) Generation-2 CNA
- 2 Cisco Nexus 5010 switches running Cisco NX-OS Release 4.1(3)N1(1)

The configuration example includes the following considerations and requirements:

1. Generation 2 CNAs that support DCBX are required.

Send documentation comments to n5kdocfeedback@cisco.com

2. Single host CNA port channel connection to a separate switch. FCoE interfaces will not be brought up if the port channel on a single switch contains more than one member port in a port channel or vPC.
3. Cisco NX-OS Release 4.1(3)N1(1) or a later release.
4. FC Features Package (FC_FEATURES_PKG) is necessary for running FCoE. If this is not installed, there will be a temporary license that will last 90 days.

This appendix includes the following sections:

- [Cisco Nexus 5000 Series Switch vPC Configuration Example, page 4-2](#)
- [Cisco Nexus 5000 Series Switch FCoE Configuration Example, page 4-5](#)

Cisco Nexus 5000 Series Switch vPC Configuration Example

This example presumes that the basic configuration has been completed on the switch (for example, IP Address (mgmt0), switchname, and password for the administrator).

This example shows how to configure the basic vPC configuration. For more information on configuring vPC, refer to the [Cisco Nexus 5000 Series vPC Quick Configuration Guide](#).



Note

The configuration must be done on both peer switches in the vPC topology.

Step 1 Enable the vPC feature on both peer switches.

```
tme-n5k-1# conf t
Enter configuration commands, one per line. End with CNTL/Z.
tme-n5k-1(config)# feature vpc
tme-n5k-1(config)#

tme-n5k-2# conf t
Enter configuration commands, one per line. End with CNTL/Z.
tme-n5k-2(config)# feature vpc
tme-n5k-2(config)#
```

Step 2 Configure the vPC domain and peer-keep alive destinations:

```
tme-n5k-1(config)# vpc domain 2
tme-n5k-1(config-vpc-domain)# peer-keepalive destination 192.165.200.229

tme-n5k-2(config)# vpc domain 2
tme-n5k-2(config-vpc-domain)# peer-keepalive destination 192.165.200.230
```



Note

In this set up, switch tme-n5k-1 has the mgmt IP address of 192.165.200.229 and switch tme-n5k-2 has the mgmt IP address of 192.165.200.230.

Step 3 Configure the port channel interface that will be used as the vPC peer-link:

```
tme-n5k-1(config)# int port-channel 1
tme-n5k-1(config-if)# vpc peer-link
```



Note

The spanning tree port type is changed to network port type on vPC peer-link. This will enable STP Bridge Assurance on vPC peer-link provided that the STP Bridge Assurance (which is enabled by default) is not disabled.

Send documentation comments to n5kdocfeedback@cisco.com

```
tme-n5k-2(config)# int port-channel 1
tme-n5k-2(config-if)# vpc peer-link
```

Step 4 Verify that the peer-keepalive can be reached:

```
tme-n5k-1(config)# show vpc peer-keepalive
vPC keep-alive status      : peer is alive
--Destination              : 172.25.182.167
--Send status              : Success
--Receive status           : Success
--Last update from peer    : (0 ) seconds, (975 ) msec
tme-n5k-1(config)#
```

```
tme-n5k-2(config)# show vpc peer-keepalive
--PC keep-alive status     : peer is alive
--Destination              : 172.25.182.166
--Send status              : Success
--Receive status           : Success
--Last update from peer    : (0 ) seconds, (10336 ) msec
tme-n5k-2(config)#
```

Step 5 Add member ports to the vpc-peer link port channel and bring up the port channel interface:

```
tme-n5k-1(config-if-range)# int po 1
tme-n5k-1(config-if)# switchport mode trunk
tme-n5k-1(config-if)# no shut
tme-n5k-1(config-if)# exit
tme-n5k-1(config)# int eth 1/39-40
tme-n5k-1(config-if-range)# switchport mode trunk
tme-n5k-1(config-if-range)# channel-group 1
tme-n5k-1(config-if-range)# no shut
tme-n5k-1(config-if-range)#
```

```
tme-n5k-2(config-if-range)# int po 1
tme-n5k-2(config-if)# switchport mode trunk
tme-n5k-2(config-if)# no shut
tme-n5k-2(config-if)# exit
tme-n5k-2(config)# int eth 1/39-40
tme-n5k-2(config-if-range)# switchport mode trunk
tme-n5k-2(config-if-range)# channel-group 1
tme-n5k-2(config-if-range)# no shut
tme-n5k-2(config-if-range)#
```

```
tme-n5k-1(config-if-range)# show int po1
port-channel 1 is up
Hardware: Port-Channel, address: 000d.ecde.a92f (bia 000d.ecde.a92f)
MTU 1500 bytes, BW 20000000 Kbit, DLY 10 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA
Port mode is trunk
full-duplex, 10 Gb/s
Beacon is turned off
Input flow-control is off, output flow-control is off
Switchport monitor is off
Members in this channel: Eth1/39, Eth1/40
Last clearing of "show interface" counters never
1 minute input rate 1848 bits/sec, 0 packets/sec
1 minute output rate 3488 bits/sec, 3 packets/sec
tme-n5k-1(config-if-range)#

tme-n5k-2(config-if-range)# show int po1
port-channell is up
Hardware: Port-Channel, address: 000d.ecdf.5fae (bia 000d.ecdf.5fae) MTU 1500 bytes,
BW 20000000 Kbit, DLY 10 usec,
```

Send documentation comments to n5kdocfeedback@cisco.com

```
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA
Port mode is trunk
full-duplex, 10 Gb/s
Beacon is turned off
Input flow-control is off, output flow-control is off
Switchport monitor is off
Members in this channel: Eth1/39, Eth1/40
Last clearing of "show interface" counters never
minute input rate 1848 bits/sec, 0 packets/sec
minute output rate 3488 bits/sec, 3 packets/sec
tme-n5k-2(config-if-range)#
```

Step 6 Create the vPC and add member interfaces:

```
tme-n5k-1(config)# int po 11
tme-n5k-1(config-if)# vpc 11
tme-n5k-1(config-if)# switchport mode trunk
tme-n5k-1(config-if)# no shut
tme-n5k-1(config-if)# int eth 1/1
tme-n5k-1(config-if)# switchport mode trunk
tme-n5k-1(config-if)# channel-group 11
tme-n5k-1(config-if)# spanning-tree port type edge trunk
tme-n5k-1(config-if)#
```



Warning

Edge port type (portfast) should only be enabled on ports connected to a single host. Connecting some devices such as hubs, concentrators, switches, or bridges to this interface when edge port type (portfast) is enabled, can cause temporary bridging loops. Caution should be used in this type of configuration

```
tme-n5k-2(config)# int po 11
tme-n5k-2(config-if)# vpc 11
tme-n5k-2(config-if)# switchport mode trunk
tme-n5k-2(config-if)# no shut
tme-n5k-2(config-if)# int eth 1/1
tme-n5k-2(config-if)# switchport mode trunk
tme-n5k-2(config-if)# channel-group 11
tme-n5k-2(config-if)# spanning-tree port type edge trunk
```



Warning

Edge port type (portfast) should only be enabled on ports connected to a single host. Connecting some devices such as hubs, concentrators, switches, or bridges to this interface when edge port type (portfast) is enabled, can cause temporary bridging loops. Caution should be used in this type of configuration.



Note

To run FCoE over a vPC topology, the port channel can only have a single member interface.



Note

The vPC number configured under the port channel interface must match on both Nexus 5000 switches. The port channel interface number does not have to match on both switches.

Step 7 Verify that the vPC interfaces are up and operational:

```
tme-n5k-1(config-if)# show vpc statistics vpc 11
port-channel111 is up
vPC Status: Up, vPC number: 11
Hardware: Port-Channel, address: 000d.ecde.a908 (bia 000d.ecde.a908)
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
```

Send documentation comments to n5kdocfeedback@cisco.com

```

reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA
Port mode is trunk
full-duplex, 10 Gb/s
Beacon is turned off
Input flow-control is off, output flow-control is off
Switchport monitor is off
Members in this channel: Eth1/1
Last clearing of "show interface" counters never
minute input rate 4968 bits/sec, 8 packets/sec
minute output rate 792 bits/sec, 1 packets/sec
tme-n5k-1(config-if)#

tme-n5k-2(config-if)# show vpc statistics vpc 11
port-channel11 is up
vPC Status: Up, vPC number: 11
Hardware: Port-Channel, address: 000d.ecdf.5fae (bia 000d.ecdf.5fae)
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA
Port mode is trunk
full-duplex, 10 Gb/s
Beacon is turned off
Input flow-control is off, output flow-control is off
Switchport monitor is off
Members in this channel: Eth1/1
Last clearing of "show interface" counters never
minute input rate 4968 bits/sec, 8 packets/sec
minute output rate 792 bits/sec, 1 packets/sec
tme-n5k-1(config-if)#

```

Cisco Nexus 5000 Series Switch FCoE Configuration Example

Once the vPC is set up between the two Nexus 5000s, we can move on to configuring the FCoE topology. This cheat sheet presumes that basic configuration has been executed on the Nexus 5000 switch that will provide IP Address (mgmt0), switchname, password for admin, etc. and that the vPC configuration has been completed as outlined in the previous section. The following steps will walk through the basic FCoE configuration necessary to set up an FCoE topology in conjunction with the vPC topology.

Step 1 Enable FCoE on the Nexus 5000:

```

tme-n5k-1(config)# feature fcoe
FC license checked out successfully
fc_plugin extracted successfully
FC plugin loaded successfully
FCoE manager enabled successfully
FC enabled on all modules successfully
tme-n5k-1(config)#

tme-n5k-2(config)# feature fcoe
FC license checked out successfully
fc_plugin extracted successfully
FC plugin loaded successfully
FCoE manager enabled successfully
FC enabled on all modules successfully
fme-n5k-2(config)#

```

Send documentation comments to n5kdocfeedback@cisco.com



Note This can take a few moments to complete.

Step 2 Create a VSAN and map it to a VLAN that has been designated to carry FCoE traffic:

```
tme-n5k-1(config)# vsan database
tme-n5k-1(config-vsan-db)# vsan 100
tme-n5k-1(config-vsan-db)# exit
tme-n5k-1(config)# vlan 100
me-n5k-1(config-vlan)# fcoe vsan 100
tme-n5k-1(config-vlan)# show vlan fcoe
VLAN      VSAN      Status
-----
100       100       Operational
tme-n5k-1(config-vlan)#

tme-n5k-2(config)# vsan database
tme-n5k-2(config-vsan-db)# vsan 101
tme-n5k-2(config-vsan-db)# exit
tme-n5k-2(config)# vlan 101
tme-n5k-2(config-vlan)# fcoe vsan 101
tme-n5k-2(config-vlan)# show vlan fcoe
VLAN      VSAN      Status
-----
101       101       Operational
tme-n5k-2(config)#
```



Note VLAN and VSAN numbers are not required to be the same.

Step 3 Configure the VLANs that are allowed to transverse the vPC links:

```
tme-n5k-1(config)# int po 11
tme-n5k-1(config-if)# switchport trunk allowed vlan 1, 100
tme-n5k-1(config-if)# show int trunk
-----
Port          Native    Status    Port
-----
Eth1/1        1         trnk-bndl Po11
Eth1/39       1         trnk-bndl Po1
Eth1/40       1         trnk-bndl Po1
Po1           1         trunking  --
Po11          1         trunking  --
-----

Port          Vlans Allowed on Trunk
-----
Eth1/1        1,100
Eth1/39       1-3967,4048-4093
Eth1/40       1-3967,4048-4093
Po1           1-3967,4048-4093
Po11          1,100
-----

Port          Vlans Err-disabled on Trunk
-----
Eth1/1        none
Eth1/39       100
Eth1/40       100
Po1           100
Po11          none
```


Send documentation comments to n5kdocfeedback@cisco.com

```

-----
Port                STP Forwarding
-----
Eth1/1              none
Eth1/39             none
Eth1/40             none
Po1                 1
Po11                1,100
tme-n5k-1(config-if)#

tme-n5k-2(config)# int po 11
tme-n5k-2(config-if)# switchport trunk allowed vlan 1, 101
tme-n5k-2(config-if)# show int trunk
-----
Port                Native      Status      Port
-----
Eth1/1              1          trnk-bndl   Po11
Eth1/39             1          trnk-bndl   Po1
Eth1/40             1          trnk-bndl   Po1
Po1                 1          trunking    --
Po11                1          trunking    --
-----

Port                Vlans Allowed on Trunk
-----
Eth1/1              1,101
Eth1/39             1-3967,4048-4093
Eth1/40             1-3967,4048-4093
Po1                 1-3967,4048-4093
Po11                1,101
-----

Port                Vlans Err-disabled on Trunk
-----
Eth1/1              none
Eth1/39             101
Eth1/40             101
Po1                 101
Po11                none
-----

Port                STP Forwarding
-----
Eth1/1              none
Eth1/39             none
Eth1/40             none
Po1                 1
Po11                1,101
tme-n5k-2(config-if)#

```

- Step 4** Create a virtual Fibre Channel interface (vfc) and add it to the VSAN that was created in the previous step:

```

tme-n5k-1(config)# int vfc 1
tme-n5k-1(config-if)# bind interface po11
Warning: VFC will not come up for pre-FIP CNA
tme-n5k-1(config-if)# no shut
tme-n5k-1(config-if)#

tme-n5k-2(config)# int vfc 1
tme-n5k-2(config-if)# bind interface po11
Warning: VFC will not come up for pre-FIP CNA
tme-n5k-2(config-if)# no shut
tme-n5k-2(config-if)#

```

Send documentation comments to n5kdocfeedback@cisco.com

```
tme-n5k-1(config)# vsan database
tme-n5k-1(config-vsan-db)# vsan 100 interface vfc 1
tme-n5k-1(config)# show vsan membership
vsan 1 interfaces:
fc2/1          fc2/2          fc2/3          fc2/4
fc2/5          fc2/6          fc2/7          fc2/8

vsan 100 interfaces:
vfc1

vsan 4079(evfp_isolated_vsan) interfaces:

vsan 4094(isolated_vsan) interfaces:
tme-n5k-1(config)#

tme-n5k-2(config)# vsan database
tme-n5k-2(config-vsan-db)# vsan 101 interface vfc 1
tme-n5k-2(config)# show vsan membership
vsan 1 interfaces:
fc2/1          fc2/2          fc2/3          fc2/4
fc2/5          fc2/6          fc2/7          fc2/8

vsan 101 interfaces:
vfc1

vsan 4079(evfp_isolated_vsan) interfaces:

vsan 4094(isolated_vsan) interfaces:
tme-n5k-2(config)#
```

Step 5 Verify that the vfc is up and operational:

```
tme-n5k-1(config-if)# show int brief
-----
Ethernet      VLAN   Type   Mode   Status   Reason           Speed
-----
Eth1/1        1      eth    trunk  up       none             10G(D)
Eth1/2        1      eth    access up       none             10G(D)
Eth1/38       1      eth    access down    SFP not inserted 10G(D)
Eth1/39       1      eth    trunk  up       none             10G(D)
Eth1/40       1      eth    trunk  up       none             10G(D)
-----

Port-channel  VLAN   Type   Mode   Status   Reason           Speed
-----
Po1           1      eth    trunk  up       none             a-10G(D)  none
Po11          1      eth    trunk  up       none             a-10G(D)  none
-----

Port   VRF      Status IP Address           Speed   MTU
-----
mgmt0  --        up      172.25.182.166       1000   1500
-----

Interface  Vsan      Admin  Admin  Status   SFP  Oper  Oper  Port
-----
vfc1       100      F      on     up       --   F    auto  --
tme-n5k-1(config-if)#

tme-n5k-2(config-if)# show int brief
-----
Ethernet      VLAN   Type   Mode   Status   Reason           Speed   Port
-----
```

Send documentation comments to n5kdocfeedback@cisco.com

```

Eth1/1      1      eth   trunk  up      none      10G(D)  11
Eth1/2      1      eth   access up      none      10G(D)  --
Eth1/38     1      eth   access down   SFP not inserted 10G(D)  --
Eth1/39     1      eth   trunk  up      none      10G(D)  1
Eth1/40     1      eth   trunk  up      none      10G(D)  1

-----
Port-channel VLAN  Type Mode  Status Reason          Speed  Protocol
-----
Po1           1      eth   trunk  up      none      a-10G(D) none
Po11          1      eth   trunk  up      none      a-10G(D) none

-----
Port  VRF          Status IP Address          Speed  MTU
-----
mgmt0 --          up      172.25.182.167      1000  1500

-----
Interface  Vsan      Admin  Admin  Status      SFP  Oper  Oper
-----
vfc1       101      F      on     up      --   F     auto  --
tme-n5k-2(config-if)#

```

Step 6 Verify that the virtual Fibre Channel interface has logged into the fabric:

```

tme-n5k-1# show flogi database
-----
INTERFACE      VSAN      FCID          PORT NAME          NODE NAME
-----
vfc1           100      0x540000 21:00:00:c0:dd:11:2a:01 20:00:00:c0:dd:11:2a:01

Total number of flogi = 1.
tme-n5k-2# show flogi database
-----
INTERFACE      VSAN      FCID          PORT NAME          NODE NAME
-----
vfc1           101      0x540000 21:00:00:c0:dd:11:2a:01 20:00:00:c0:dd:11:2a:01

Total number of flogi = 1.

```

Step 7 Verify that the vPC is up and operational:

```

tme-n5k-1(config-if)# show vpc statistics vpc 11
port-channel11 is up
vPC Status: Up, vPC number: 11
Hardware: Port-Channel, address: 000d.ecde.a908 (bia 000d.ecde.a908)
  MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
  reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA
  Port mode is trunk
  full-duplex, 10 Gb/s
  Beacon is turned off
  Input flow-control is off, output flow-control is off
  Switchport monitor is off
  Members in this channel: Eth1/1
  Last clearing of "show interface" counters never
  1 minute input rate 4968 bits/sec, 8 packets/sec
  1 minute output rate 792 bits/sec, 1 packets/sec

tme-n5k-2(config-if)# show vpc statistics vpc 11
port-channel11 is up
vPC Status: Up, vPC number: 11

```

Send documentation comments to n5kdocfeedback@cisco.com

```
Hardware: Port-Channel, address: 000d.ecdf.5fae (bia 000d.ecdf.5fae)
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA
Port mode is trunk
full-duplex, 10 Gb/s
Beacon is turned off
Input flow-control is off, output flow-control is off
Switchport monitor is off
Members in this channel: Eth1/1
Last clearing of "show interface" counters never
1 minute input rate 4968 bits/sec, 8 packets/sec
1 minute output rate 792 bits/sec, 1 packets/sec
```



CHAPTER 5

FCoE with Cisco Nexus 4000 Series Switch Configuration Example

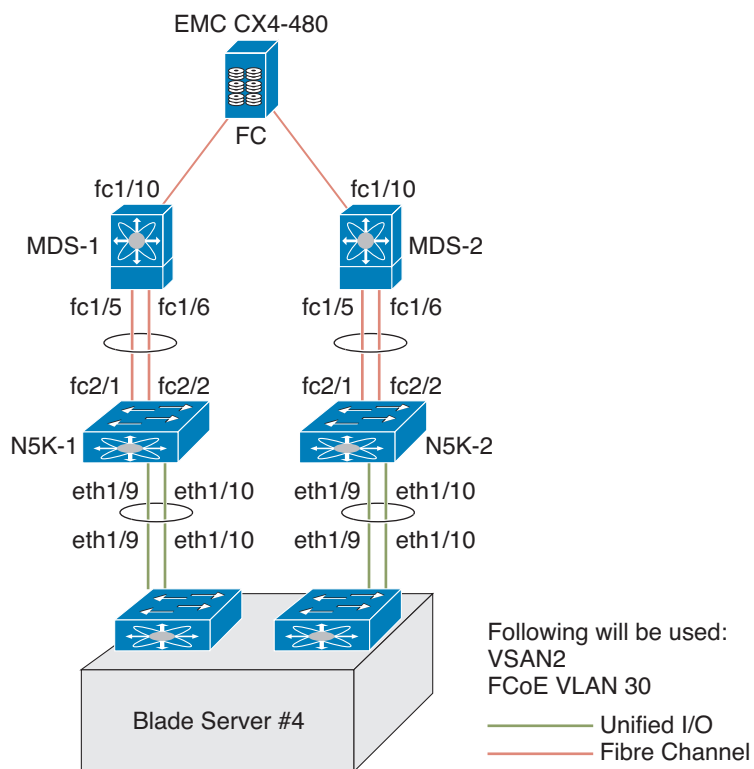
This section includes a configuration example on how to configure an IBM blade server connecting to a Cisco Nexus 4000 Series switch which is then connected to a Cisco Nexus 5000 Series switch which accesses FC storage on a Cisco MDS 9000 Series Family switch using FCoE. Because the Cisco Nexus 4000 Series switch is a FIP snooping bridge, the FLOGI done by the CNAs do not login on the Cisco Nexus 4000 Series switch but onto the Cisco Nexus 5000 Series switch, which is the FCF. Creation of the vFC interface for the Cisco Nexus 4000 Series switch blade servers does not change whether the Cisco Nexus 5000 Series switch is in switching or NPV mode. Where the actual fabric login happens is determined by the mode on the Cisco Nexus 5000 Series switch.

- Cisco Nexus 5000 Series switch in switching mode—Login is on the Cisco Nexus 5000 Series switch.
- Cisco Nexus 5000 Series switch in NPV mode—Login will be on the Cisco MDS 9000 Series Family switch or any FC switch upstream with NPIV configured.

In this example, the Cisco Nexus 5000 Series switch is in switching mode. [Figure 5-1](#) shows the topology used in the example.

Send documentation comments to n5kdocfeedback@cisco.com

Figure 5-1 Nexus 4000 FCoE Lab Topology



The following hardware was used:

- IBM Blade Chassis model BCH
- IBM HS22 blade server running Windows 2003 using the Qlogic QMI8142
- Cisco Nexus 4000 Series switch running Cisco NX-OS Release 4.1(2)E1(1)
- Cisco Nexus 5010 switch running Cisco NX-OS Release 4.1(3)N1(1)
- Cisco MDS 9124 Director switch running Cisco SAN-OS Release 4.1(3a)
- EMC CX4-480

This appendix includes the following sections:

- [Cisco Nexus 5000 Series Switch in Switching Mode, page 5-3](#)
- [Configuring a SAN Port Channel on the Cisco Nexus 5000 Series Switch to the Cisco MDS Directory Series, page 5-4](#)
- [Configuring a Port Channel on a Cisco Nexus 5000 Series Switch to a Cisco Nexus 4000 Series Switch, page 5-5](#)
- [Configuring a Virtual Fibre Channel Interface on a Cisco Nexus 4000 Series Switch, page 5-6](#)
- [Configuring a VSAN on the Cisco Nexus 5000 Series Switch, page 5-6](#)

Send documentation comments to n5kdocfeedback@cisco.com

Cisco Nexus 5000 Series Switch in Switching Mode

Before following the steps in this example, be sure to complete a basic configuration on the Cisco Nexus 5000 Series switch (for example, IP Address (mgmt0), switch name, and password for the administrator) and FCoE has not been enabled.

To use this configuration example in production, you must have the FC Features Package license installed otherwise there will be a temporary license that expires after 90 days. When the license expires, the feature is disabled.

On the Cisco Nexus 5000 Series switch, by default FCoE is not enabled.

This example shows how to enable FCoE:

```
n5k-2# show interface brief
-----
Ethernet VLAN Type Mode Status Reason Speed Port
Interface                               Ch #
Eth1/1 1 eth access up none 10G(D) --
Eth1/2 1 eth access up none 10G(D) --
[snip]
Eth2/4 1 eth access down SFP not inserted 10G(D) --
-----
Port VRF Status IP Address Speed MTU
-----
mgmt0 -- up 172.25.182.164 1000 1500
```



Note

There are no FC interfaces, even though there is a 4x4 GEM card installed in the Cisco Nexus 5010 switch.

```
n5k-2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
n5k-2(config)# feature fcoe
FC license checked out successfully fc_plugin extracted successfully FC plugin loaded
successfully FCoE manager enabled successfully FC enabled on all modules successfully
```



Note

Beginning with Cisco NX-OS Release 4.1(3)N1(1), the switch does not need to be reboot when you enable FCoE. The Cisco Nexus 5000 Series switch is in switching mode by default when FCoE is enabled.

```
n5k-2(config)# show feature
Feature Name Instance State
fcsp 1 disabled
fcoe 1 enabled
fex 1 enabled

n5k-2(config)# show interface brief
-----
Interface Vsan Admin Admin Status SFP Oper Oper Port
Mode Trunk Mode Speed Channel
Mode (Gbps)
-----
fc2/1 1 auto on down sw1 -- --
fc2/2 1 auto on down sw1 -- --
fc2/3 1 auto on down sw1 -- --
fc2/4 1 auto on sfpAbsent -- -- --
```

Send documentation comments to n5kdocfeedback@cisco.com

```
-----
Ethernet VLAN Type Mode Status Reason Speed Port
Interface Ch #
-----
```

```
Eth1/1 1 eth access up none 10G(D) --
Eth1/2 1 eth access up none 10G(D) --
```



Note

Use the **show interface brief** command to show the FC interfaces.

Configuring a SAN Port Channel on the Cisco Nexus 5000 Series Switch to the Cisco MDS Directory Series

This example shows how to configure a SAN port channel on the Cisco Nexus 5000 Series switch that is connected to a Cisco MDS 9000 Director. For redundancy, Cisco recommends that you create a SAN port channel from the FC interfaces.

Step 1 Configure a SAN port channel on the Cisco Nexus 5000 Series switch.

```
Fn5k-2# configure terminal
n5k-2(config)# interface san-port-channel 1
n5k-2(config-if)# interface fc2/1-2
n5k-2(config-if)# channel-group 1
```



Note

After you add fc2/1 fc2/2 to san-port-channel 1 you need to disable the port channel. This must also be done on the switch at the other end of the port channel. Then, shut the interfaces at both ends to bring them up.

```
n5k-2(config-if)# no shut
n5k-2(config-if)# interface san-port-channel 1
n5k-2(config-if)# no shut
n5k-2(config-if)# show san-port-channel database
san-port-channel 1
Administrative channel mode is on Operational channel mode is on Last membership
update is successful 2 ports in total, 0 ports up Age of the port-channel is
0d:00h:17m:14s
Ports: fc2/1 [down] fc2/2 [down]
n5k-2(config-if)#
```



Note

The SAN port channel is currently down because the Cisco MDS 9000 Series Director has not been configured.

Step 2 Configure the Cisco MDS 9124 switch to create a port channel between the Cisco Nexus 5000 Series switch and the Cisco MDS 9124 switch.



Note

With the SAN port channel on the Cisco Nexus 5000 configured to the MDS, you will need to perform the same configuration on the Cisco MDS 9000 Series switch. A SAN port channel configuration on the Cisco MDS 9000 Series switch is called a port channel.

```
mds9124-2# configure terminal
```


Send documentation comments to n5kdocfeedback@cisco.com

```
mds9124-2(config)# interface port-channel 1
mds9124-2(config-if)# interface fc1/5, fc1/6
mds9124-2(config-if)# channel-group 1 force
```



Note

After you add fc1/5 fc1/6 to port-channel 1 you need to disable the port channel. This must also be done on the switch at the other end of the port channel. Then, shut the interfaces at both ends to bring them up.



Note

```
mds9124-2(config-if)# no shut
mds9124-2(config-if)# interface port-channel 1
mds9124-2(config-if)# no shut
```

- Step 3** Verify that the SAN port channel on the Cisco Nexus 5000 Series switch is up and running. Use the **show san-port-channel database** command to show the SAN port channel configuration.

```
n5k-2(config-if)# show san-port-channel database
san-port-channel 1
Administrative channel mode is on
Operational channel mode is on
Last membership update is successful
2 ports in total, 2 ports up
First operational port is fc2/2
Age of the port-channel is 0d:00h:25m:10s
Ports: fc2/1 [up]
fc2/2 [up] *
```

Configuring a Port Channel on a Cisco Nexus 5000 Series Switch to a Cisco Nexus 4000 Series Switch

This example shows how to configure a port channel on the Cisco Nexus 5000 Series switch that is connected to the Cisco Nexus 4000 Series switch.

- Step 1** Configure the port channel on the Cisco Nexus 5000 Series switch.

The port channel is configured to provide redundancy for traffic coming from the Cisco Nexus 4000 Series switch to the Cisco Nexus 5000 Series switch. In this example, all VLANs can traverse the port channel. The FCoE VLAN and the native VLAN must be allowed to traverse the port channel. In production environments, Network Administrators may designate other VLANs to traverse this network.

```
n5k-2# configure terminal
n5k-2(config)# feature lacp
n5k-2(config)# interface port-channel 2 mode active
n5k-2(config-if)# interface eth1/9-10
n5k-2(config-if)# channel-group 2
n5k-2(config)# interface port-channel 2
n5k-2(config-if)# switchport mode trunk
n5k-2(config-if)# no shut
n5k-2#
```

- Step 2** Configure the port channel on the Cisco Nexus 4000 Series switch.

```
bch1-n4k-b9# configure terminal
bch1-n4k-b9(config)# feature lacp
bch1-n4k-b9(config)# interface port-channel 20
```

Send documentation comments to n5kdocfeedback@cisco.com

```
bch1-n4k-b9(config-if)# interface eth1/15-16
bch1-n4k-b9(config-if)# channel-group 2 mode active
bch1-n4k-b9(config)# interface port-channel 2
bch1-n4k-b9(config-if)# switchport mode trunk
bch1-n4k-b9(config-if)# no shut
bch1-n4k-b9(config-if)#
```

Configuring a Virtual Fibre Channel Interface on a Cisco Nexus 4000 Series Switch

This example shows how to configure a vFC interface on a Cisco Nexus 4000 Series switch.

-
- Step 1** On the Cisco Nexus 5000 Series switch, configure a VSAN to match the production VSAN on the Cisco MDS 9000 Series switch. This is a one-time configuration.
 - Step 2** On the Cisco Nexus 5000 Series switch, configure an FCoE VLAN to map to the VSAN (VLAN-to-VSAN mapping). This is one-time configuration.
 - Step 3** On the Cisco Nexus 4000 Series switch, configure a FIP snooping VLAN that matches the FCoE VLAN on the Nexus 5000 Series switch. This is a one-time configuration.
 - Step 4** On the Cisco Nexus 4000 Series switch, configure the uplinks to allow FCoE traffic (FIP snooping).
 - Step 5** On the Cisco Nexus 4000 Series switch blade server, configure the Ethernet interfaces for FCoE traffic.
 - Step 6** On the Cisco Nexus 5000 Series switch, configure the vFCs.
 - Step 7** On the Cisco Nexus 4000 Series switch blade server, bind the vFC to the MAC address of the blade server.
 - Step 8** Verify that the vFC is in the correct VSAN.



Note Completing the above tasks ensure that the connection to an FCoE CNA on the blade server from the Nexus 4000 is successful.

Configuring a VSAN on the Cisco Nexus 5000 Series Switch

You can configure a VSAN on the Cisco Nexus 5000 Series switch using Fabric Manager, Device Manager, or the CLI. This example shows CLI configuration tasks and Fabric Manager or Device Manager GUI tasks.

This example shows the storage on the Cisco MDS 9000 Series resides on VSAN 2. Configure the VSAN to ensure that the vFCs configured on the Cisco Nexus 5000 Series switch can communicate with the storage device.

```
n5k-2# configure terminal
n5k-2(config)# vsan database
n5k-2(config-vsan-db)# vsan 2
n5k-2(config-vsan-db)# show vsan vsan 1 information
name:VSAN0001 state:active
interoperability mode:default
loadbalancing:src-id/dst-id/oxid
operational state:up
vsan 2 information
```

Send documentation comments to n5kdocfeedback@cisco.com

```
name:VSAN0002 state:active
interoperability mode:default
loadbalancing:src-id/dst-id/oxid
operational state:down
  vsan 4079:evfp_isolated_vsan
  vsan 4094:isolated_vsan
```

Configuring An FCoE VLAN on the Cisco Nexus 5000 Series Switch

You can configure a VLAN and then map the VLAN to a particular VSAN using the CLI. Fabric Manager and Device Manager can not be used for this configuration. Cisco recommends that you configure a separate VLAN for FCoE traffic and separate VLANs for standard Ethernet traffic.

This example shows how to create the FCoE VLAN:

```
n5k-2# configure terminal
n5k-2(config)# vlan 30
n5k-2(config-vlan)# fcoe vsan 2
n5k-2(config-vlan)# show vlan fcoe

VLAN
VSAN
Status
-----
-----
-----
30
2
Operational
```

Configuring a FIP Snooping VLAN on the Cisco Nexus 4000 Series Switch

On the Cisco Nexus 4000 Series switch, by default the FIP snooping feature is disabled. Cisco recommends that during the basic configuration, when prompted, you should enable FCoE and FIP snooping and configure, for example, the appropriate Class of Service (CoS) no drop, MTU, and QoS, without having to manually configure these features after the initial configuration.

The example shows how to verify that FIP snooping is enabled:

```
bch1-n4k-b9# show feature
Feature Name Instance State
tacacs 1 disabled lacp 1 enabled [snip] fipsm 1 enabled
```

With the FCoE VLAN configured on the Cisco Nexus 5000 Series switch as VLAN 30, then the same VLAN number must be used to create the VLAN on the Cisco Nexus 4000 Series switch and the VLAN must be configured as a FIP snooping VLAN.

This example shows how to configure the VLAN on the Cisco Nexus 4000 Series switch:

```
bch1-n4k-b9# configure terminal
bch1-n4k-b9(config)# vlan 30
bch1-n4k-b9(config-vlan)# fip-snooping enable
```

Send documentation comments to n5kdocfeedback@cisco.com

Configuring the Cisco Nexus 4000 Series Switch Uplinks To Allow FCoE Traffic

In this example, we have already created the port channel that allows all VLANs to traverse the uplink between the Cisco Nexus 4000 Series switch and the Cisco Nexus 5000 Series switch from the previous section. The uplink (in this case a port channel) must be enabled to do FIP snooping with a port type mode of fcf.

This example shows how to configure the uplink:

```
bch1-n4k-b9# configure terminal
bch1-n4k-b9(config)# interface port-channel 20
bch1-n4k-b9(config-if)# fip-snooping port-mode fcf
```

Configuring Blade Server Ethernet Interfaces on the Cisco Nexus 4000 Series Switch For FCoE Traffic

You can configure the blade server using the CLI. Fabric Manager and Device Manager can not be used for this configuration.

Ensure that the FCoE VLAN (VLAN 30) can traverse the Ethernet interface on the blade server (Ethernet 1/4). In most cases, the CNA ports allow for both regular Ethernet traffic and FCoE traffic that resides on different VLANs. By default, all Ethernet interfaces on the Cisco Nexus 4000 Series switch is in access mode and resides on VLAN 1.

This example shows how to configure the Ethernet interface to allow multiple VLANs (trunk):

```
bch1-n4k-b9#configure terminal
bch1-n4k-b9(config)#interface ethernet 1/4
bch1-n4k-b9(config-if)# switchport mode trunk
bch1-n4k-b9(config-if)# switchport trunk allowed vlan 1,30
```



Note

The above command is not needed but if you want to specify the allowed VLANs, make sure the FCoE VLAN is on the allowed list as shown in the example.

```
bch1-n4k-b9(config-if)# spanning-tree port type edge trunk
Warning: Edge port type (portfast) should only be enabled on ports connected to a single
host. Connecting hubs, concentrators, switches, bridges, etc... to this interface when
edge port type (portfast) is enabled, can cause temporary bridging loops.
Use with CAUTION
```

Creating vFC Interfaces on the Nexus 5000 - CLI

When the trunk configuration is complete, create the vFC interface on the Cisco Nexus 5000 Series switch. You can use Device Manager or the CLI to configure the vFC interface.

Because the CNA is connected on Ethernet interface eth1/4 on the Cisco Nexus 4000 Series switch and is not physically connected to the Cisco Nexus 5000 Series switch, you must bind the vFC to the MAC address of the CNA that is doing FCoE. At this time, Qlogic is the only vendor that does FCoE on the blade server that is interoperable with the Cisco Nexus 4000 Series switch. Qlogic provides 2 separate MAC addresses, one for the standard Ethernet traffic and another specifically for FCoE.

Send documentation comments to n5kdocfeedback@cisco.com

This example shows how to identify the MAC address from the specific blade server in the IBM blade chassis.

```
bch1-n4k-b9# show fip-snooping vlan-discovery
Legend:
Interface VLAN FIP MAC
Eth1/4 1 00:c0:dd:04:0c:df
Eth1/5 1 00:c0:dd:04:0d:13
```

Use the MAC address that has been identified on the blade server to create the vFC for this blade server on the Cisco Nexus 5000 series switch.

This example shows that the vFC is moved into VSAN 2. As a best practice in creating the vFC number to devices on the Cisco Nexus 4000 Series switch, you should create a numbering scheme that can easily identify where the vFCs are mapped to which blade server on which blade chassis. For this example, we are using the blade server in slot 4 on the first IBM blade chassis, which we have named BCH1. In this example, the vFC for this blade server is interface vfc104.

```
n5k-2# configure terminal
n5k-2(config)# interface vfc 104
n5k-2(config-if)# bind mac-address 00:c0:dd:04:0c:df
n5k-2(config-if)# no shutdown
n5k-2(config-if)# show vsan membership
vsan 1 interfaces:
fc2/1 fc2/2 fc2/3 fc2/4 san-port-channel 1 vfc104
vsan 2 interfaces:
vsan 4079(evfp_isolated_vsan) interfaces:
vsan 4094(isolated_vsan) interfaces:
n5k-2(config-if)# vsan database 0
this will get to the VSAN database
n5k-2(config-vsan-db)# vsan 2 interface vfc104
n5k-2(config-vsan-db)# show vsan membership
vsan 1 interfaces: fc2/1 fc2/2 fc2/3 fc2/4 san-port-channel 1
vsan 2 interfaces:
vfc104
vsan 4079(evfp_isolated_vsan) interfaces:
n5k-2# show interface vfc104
vfc104 is up
```

Bound MAC is 00:c0:dd:04:0c:df FCF priority is 128 Hardware is Virtual Fibre Channel Port WWN is 20:67:00:0d:ec:b2:b9:bf Admin port mode is F, trunk mode is on snmp link state traps are enabled Port mode is F, FCID is 0xcd0000 Port vsan is 2 [snip]

Configuring The vFC Interface Using Device Manager

This example shows how to use Device Manager to create the vFC interface.

Send documentation comments to n5kdocfeedback@cisco.com


- Step 1** Open Device Manager and login to the Cisco Nexus 5000 Series switch.

Figure 5-2 Device Manager Login Window




- Step 2** From the Device Manager menu, choose Interface > Virtual Interfaces > Fibre Channel to configure one vFC. You can also use the Quick Configuration Tool to configure multiple vFCs and bind them to physical interfaces at one time.

Figure 5-3 Device Manager Menu



- Step 3** From the Virtual FC Interfaces window, click Create to create the vFC.

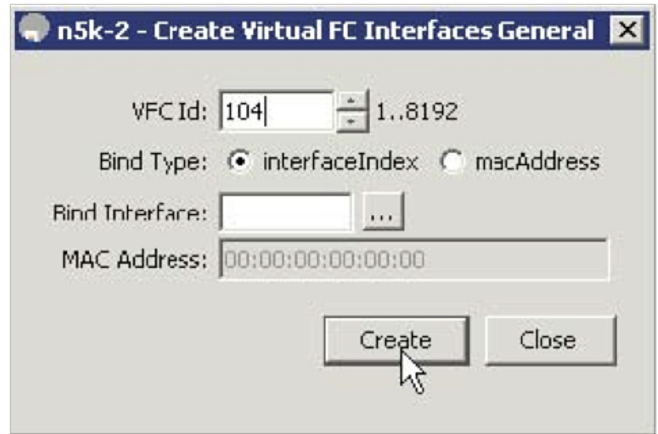
Figure 5-4 Virtual FC Interfaces Window



- Step 4** In the Create Virtual FC Interfaces General window, enter the VFC Id, Bind Type and the interface (physical or MAC address depending on the bind type) and click Create. The window is redisplayed showing the vFCs with the new vFC ID.

Send documentation comments to n5kdocfeedback@cisco.com

Figure 5-5 Create Virtual FC Interfaces General Window




Note

As a best practice, create a vFC that is recognizable of the vFC back to the blade server. For example, 104 correlates to BCH1 on blade server 4.

- Step 5** From the Virtual FC Interfaces window, choose Bind Type > macAddress.


Figure 5-6 Changing The Bind Type From Interface to Mac Address



Once the Bind Type is set to macAddress, you can enter the MAC address for the blade server in the Bind MAC Address column. In this example, 00:c0:dd:04:0c:df is the MAC address. By default, the VSAN membership is set down and VSAN 1. You can edit these sections for example, VSAN 2 and up).

- Step 6** Click on Apply to commit the changes and then click Refresh to validate the vFC is up.

Figure 5-7 The Configured vFC MAC Address in Device Manager



This completes the configuration of FCoE on the Cisco Nexus 4000 Series switch uplinked to the Cisco Nexus 5000 Series switch. The fabric management, for example, zoning and LUN masking, is managed with the existing SAN administrators tools. The vFC appear in Fabric Manager as a normal FC device but instead of a solid line to the host, a dash line is shown from the Cisco Nexus 5000 Series switch to the host.

Send documentation comments to n5kdocfeedback@cisco.com

Figure 5-8 Fabric Manager View With FCoE Devices

The screenshot shows the Fabric Manager 4.2(1a) interface. The left pane shows the Logical Domain tree with the following structure:

- DataCenter
 - SAN
 - Fabric_md9124-1
 - Fabric_md9124-2
 - All VSANs
 - VSAN0001
 - End Devices
 - VSAN Attributes
 - Domain Manager
 - Port Security
 - Fabric Binding
 - FICON
 - FSPF
 - SNF
 - Advanced
 - Groups

The center pane shows the Information table for the selected device:

| Name | Domain Id | VSAN WWN | Principal | Status | Model | Release | UpTime |
|---------------------|-------------------------------|-------------------------------|-------------------------------|--------|---------------------------|------------------|-------------------|
| md9-2 | Dxcl(205) | Cisco 20:02:00:0d:ec:b2:b9:81 | Cisco 20:02:00:0d:ec:b0:d0:81 | OK | MSK-C1010P-BF 4.1(3)N1(1) | 4 days, 13:08:24 | |
| md9124-2 Outsc(100) | Cisco 20:02:00:0d:ec:b0:d0:81 | self | | ok | DS-C9-24 | 4.1(3A) | 15 days, 22:55:18 |

The right pane shows a network diagram for Fabric_md9124-2. The diagram includes the following components:

 - PE2900-E3/11 (connected to md9-2)
 - md9-2 (connected to PE2900-E3/11)
 - Clarian b4401#5 (connected to md9-2)
 - md9124-2 (connected to md9-2)
 - md9124-2 fe1/1 (connected to md9124-2)
 - md9124-2 fe1/2 (connected to md9124-2)

The bottom status bar shows: 2 rows | admin | 6:50:36 PM | 196 of 304 | Install: Standalone

2682241



CHAPTER 6

Using FCoE NPV

This chapter describes how to configure Fiber Channel over Ethernet (FCoE) N-port Virtualization (NPV) on the Cisco Nexus 5000 Series devices.

This chapter includes the following sections:

- [Information About FCoE NPV, page 6-1](#)
- [FCoE NPV Licensing, page 6-1](#)
- [VNP Ports, page 6-2](#)
- [FCoE NPV Configuration, page 6-2](#)
- [FCoE and Enhanced vPC Considerations, page 6-6](#)
- [SAN Boot of Initiators via LACP-Based Host vPC, page 6-7](#)
- [FCoE Functionality Using the Adapter-FEX, page 6-9](#)

Information About FCoE NPV

Beginning with Cisco NX-OS Release 5.0(3)N2(1) and later releases, FCoE NPV is supported on the Cisco Nexus 5000 Series devices. The FCoE NPV feature is an enhanced form of FCoE Initialization Protocol (FIP) snooping that provides a secure method to connect FCoE-capable hosts to an FCoE-capable FCoE forwarder (FCF) device. The FCoE NPV feature provides the following benefits:

- FCoE NPV does not have the management and troubleshooting issues that are inherent to managing hosts remotely at the FCF.
- FCoE NPV implements FIP snooping as an extension to the NPV function while retaining the traffic-engineering, vsan-management, administration, and trouble shooting aspects of NPV.
- FCoE NPV and NPV together allow communication through FC and FCoE ports at the same time, which provides a smooth transition when moving from FC to FCoE topologies.

FCoE NPV Licensing

You can enable FCoE NPV by choosing one of the following methods:

- Enable FCoE and then enable NPV—This method requires that you enable FCoE first using the **feature fcoe** command and then you enable NPV by using the **feature npv** command. When FCoE is enabled, the default mode of operation is FC switching. When you enable NPV, the mode changes to NPV mode. Switching to NPV mode automatically performs a write erase and reloads the system.

Send documentation comments to n5kdocfeedback@cisco.com

After the reload, the system comes up in NPV mode. To exit NPV mode and return to FC switching mode, enter the **no feature npv** command. Exiting NPV mode also triggers a write erase and a device reload. This method requires the Storage Protocols Services Package (FC_FEATURES_PKG) license.

- Enable FCoE NPV—When you enable FCoE NPV using the **feature fcoe-npv** command, the mode changes to NPV. When you use this method, a write erase and reload does not occur. This method requires a separate license package (FCOE_NPV_PKG). This license is also included in the Storage Protocol Services License.



Note

The **feature fcoe-npv** command is only for FCoE (not for FC and does not allow conversion of ethernet ports to FC ports). The **feature fcoe** command is for FC and FCoE. The **feature npv** command is for node port virtualization.

VNP Ports

Connectivity from an FCoE NPV bridge to the FCF is supported only over point-to-point links. These links can be individual Ethernet interfaces or port channel interfaces. For each FCF connected to an Ethernet interface, a vFC interface must be created and bound to it. These vFC interfaces must be configured as VNP ports.

On the VNP port, the FCoE NPV bridge emulates an FCoE-capable host with multiple enodes, each with a unique enode MAC address. By default, the VNP port is enabled in trunk mode.

Multiple VSANs can be configured on the VNP port. The FCoE VLANs that correspond to the VNP port VSANs must be configured on the bound Ethernet interface.



Note

VNP ports on the Cisco Nexus 5000 Series device emulate an FCoE capable host with multiple Ethernet nodes, each with unique Fabric Provided MAC-Addresses (FPMA).

FCoE NPV Configuration

As seen in [Figure 6-1](#), the FCoE-NPV device is proxying FIP control messages and the Fabric Login (FLOGI) between the converged network adapter (CNA) and the FCoE FCF device. An FCoE NPV, unlike a FIP snooping bridge, is VSAN-aware and will take VSANs into account when mapping (or pinning) logins from the CNA to an FCF uplink. FLOGI from the initiators (eNodes), are load balanced between the two links of each port channel interface (VNP) connecting the FCoE NPV and FCF devices.



Note

In Cisco Nexus 5000 Series devices, the FCoE NPV feature does not convert FLOGI to FDISC.

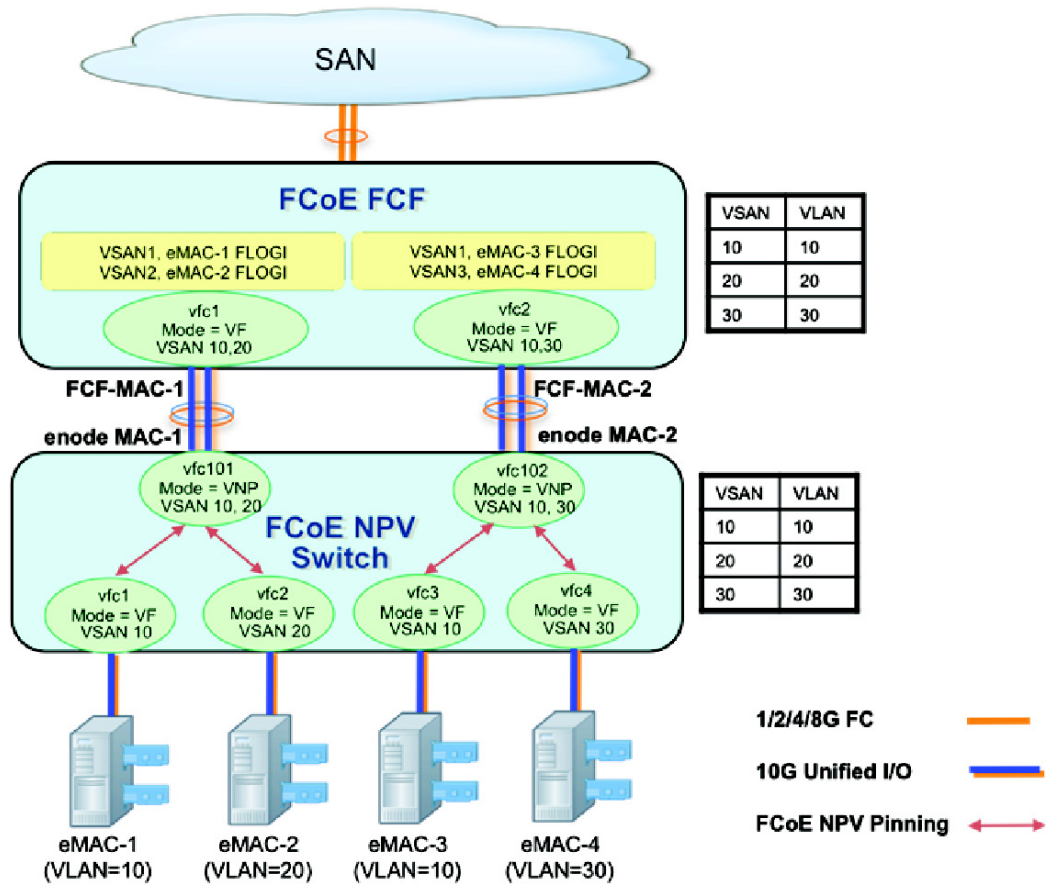


Note

Spanning Tree Protocol (STP) is disabled for FCoE VLANs on Ethernet interfaces that are bound to VNP, VF, and VE interfaces.

Send documentation comments to n5kdocfeedback@cisco.com

Figure 6-1 FCoE NPV Pinning, FLOGI-Based on VSANs



On an FCoE NPV device, the VFC interfaces bound to server-facing ports are configured in VF mode and the VFC interfaces facing the FCoE FCF are configured as VNP ports. In the following configuration example, the FCoE NPV device uses the FCOE_NPV_PKG license to enable the FCoE NPV feature. We recommend that you selectively allow a specific list of VSANs instead of allowing all VSANs on the VNP port.

```

switch(config)# feature fcoe-npv
switch(config)# vsan database
switch(config-vsan-db)# vsan 1-2
switch(config)# vlan 10
switch(config-vlan)# fcoe vsan 1
switch(config)# vlan 20
switch(config-vlan)# fcoe vsan 2
switch(config)# interface eth1/1
switch(config-if)# switchport mode trunk
switch(config)# interface Eth1/2
switch(config-if)# switchport mode trunk
switch(config)# interface vfc1
switch(config-if)# bind interface eth1/1
switch(config-if)# switchport trunk allowed vsan 1
switch(config-if)# no shut
switch(config)# interface vfc2
switch(config-if)# bind interface eth1/2
switch(config-if)# switchport trunk allowed vsan 2
switch(config-if)# no shut
switch(config)# interface vfc101

```

Send documentation comments to n5kdocfeedback@cisco.com

```
switch(config-if)# switchport trunk allowed vsan 10
switch(config-if)# switchport trunk allowed vsan add 20
switch(config-if)# no shut
switch(config)# vsan database
switch(config-vsan-db)# vsan 1 interface vfc1, vfc101
switch(config-vsan-db)# vsan 2 interface vfc2
```

QoS Requirements

For Cisco Nexus 5500 Platform devices, you must configure class-fcoe in the QoS policy maps of all types.

This example shows how to configure class-fcoe in all QoS policy maps:

```
switch# config t
switch(config)# system qos
switch(config-sys-qos)# service-policy type qos input fcoe-default-in-policy
switch(config-sys-qos)# service-policy type queuing input fcoe-default-in-policy
switch(config-sys-qos)# service-policy type queuing output fcoe-default-out-policy
switch(config-sys-qos)# service-policy type network-qos fcoe-default-nq-policy
```



Note

The above QoS configuration is essential for Cisco Nexus 5548 and Cisco Nexus 5596 platforms running Cisco NX-OS Release 5.0(3)N2(1) or earlier. Starting with NX-OS Release 5.1(3)N1(1), this configuration is applied automatically unless a user defined custom policy for FCoE traffic already exists. If a user defined policy is already configured, the default QoS policy is not applied.

FCoE NPV Features

The following FiberChannel NPV features are as follows:

- Automatic traffic mapping
- Static traffic mapping
- Disruptive load balancing
- FCoE forwarding in the FCoE NPV bridge
- CoE frames received over VNP ports are forwarded only if the L2_DA matches one of the FCoE MAC addresses assigned to hosts on the VF ports. Otherwise, they are discarded.



Note

FCoE NPV over port channel VNP ports use automatic traffic mapping only for FIP negotiations. FCoE traffic distribution over port channel VNP ports is based on the computed hash value.

Interoperability with FCoE-Capable Devices

Beginning with Cisco NX-OS Release 5.0(3)N2(1), the Cisco Nexus 5000 Series device interoperates with the following FCoE-capable devices:

- Cisco MDS 9000 Series Multilayer devices enabled to perform FCF functions (FCoE NPV and VE).
- Cisco Nexus 7000 Series devices enabled to perform FCF functions (FCoE NPV and VE).
- Cisco Nexus 4000 Series devices enabled for FIP Snooping.

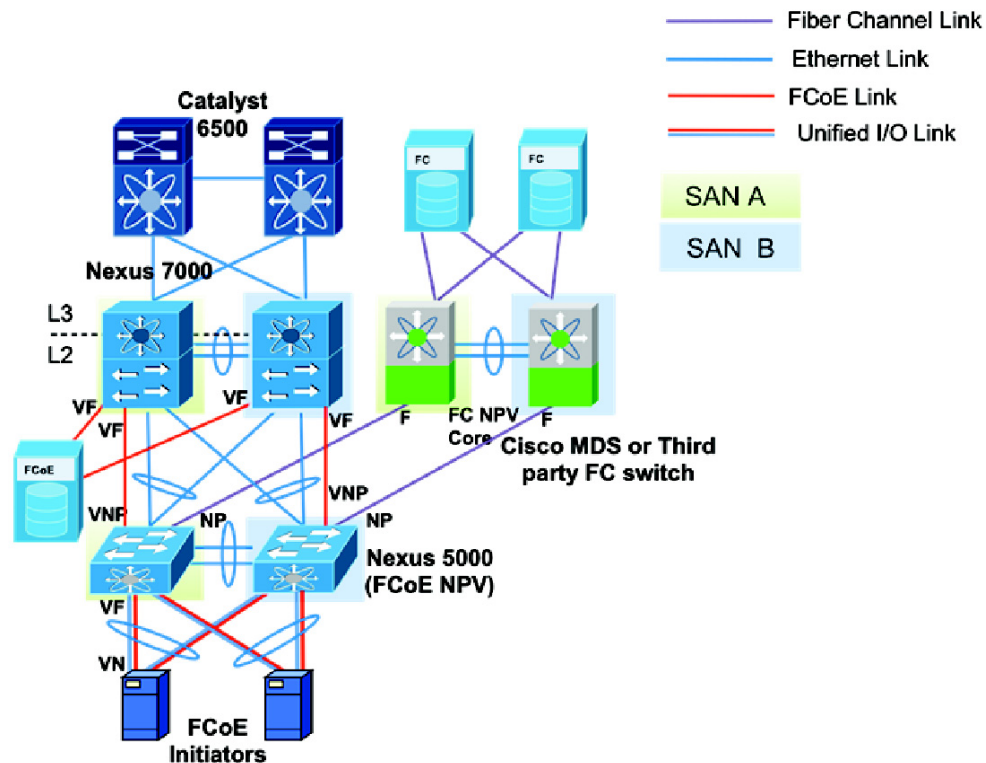
Send documentation comments to n5kdocfeedback@cisco.com

Network Design using Cisco Nexus 5000, Cisco Nexus 7000, and Cisco MDS 9500 Devices

Beginning with Cisco NX-OS Release 5.2, the Cisco Nexus 7000 Series devices supports FCoE functionality with a 32-port 10G SFP+ F1 linecard. On the Cisco MDS 9500 Multilayer Director, FCoE is supported on the 10-Gbps 8-port FCoE Module from Cisco NX-OS Release 5.2 and later releases. Based on the SAN design requirements, a combination of Cisco Nexus 7000 Series, Cisco Nexus 5000 Series, and Cisco MDS 9000 Series devices can be used to design single-hop, multi hop, or multi tier FCoE networks.

Figure 6-2 shows one recommended multi tier FCoE design topology using Cisco Nexus 7000 Series, Cisco Nexus 5000 Series, and Cisco MDS 9000 Series devices. Cisco Nexus 5000 Series devices can run FCoE NPV and FC NPV simultaneously, providing server-CNA connectivity to FCF devices while maintaining connectivity with existing SAN networks.

Figure 6-2 Converged Multi-hop FCoE Network Design Using FCoE NPV



331779

Custom QoS Settings for ETS

By default, the Cisco Nexus 5000 Series, Cisco Nexus 7000 Series, and Cisco MDS 9000 Series devices use a CoS value of 3 for the FCoE traffic. When FCoE is enabled on a Cisco Nexus 5000 Series device, CoS 3 is automatically configured for a no-drop service (PFC setting), and in the case of congestion (ETS setting) a guaranteed bandwidth of 50 percent is allocated to FCoE traffic.

The default ETS setting for Cisco Nexus 7000 Series devices and the Cisco MDS 9000 Series devices is 70 percent for FCoE no-drop traffic class.

The following commands are required on the Cisco Nexus 7000 Series device to configure QoS settings for the FCoE traffic:

Send documentation comments to n5kdocfeedback@cisco.com

```
N7K-1# configure terminal
N7K-1(config)# system qos
N7K-1(config-sys-qos)# service-policy type network-qos default-nq-7e-policy
```



Note

For Cisco Nexus 7000 Series, Cisco Nexus 5000 Series, and Cisco MDS 9000 Series devices, you should leave the default CoS value of 3 for the FCoE traffic while ensuring the same end-to-end guaranteed bandwidth gets allocated for this FCoE no-drop class.

Use of Consolidated Versus Dedicated Links

The cost/benefit tradeoff for consolidated links and dedicated links for FCoE is discussed in detail in the “Consolidated Links and Dedicated Links for FCoE” section, which covers the design considerations and link type selection you should consider when designing FCoE networks using Cisco Nexus 7000 Series and Cisco MDS 9000 Series devices.

We recommend that you use dedicated links to carry the FCoE traffic between Cisco Nexus 7000 Series and Cisco Nexus 5000 Series devices because FCoE links are part of storage Virtual Device Contents (VDCs) and not part of Ethernet VDCs. Consolidated links go to the storage VDCs of the Cisco Nexus 7000 Series devices drop Ethernet traffic carried over the link. Similarly consolidated links that are configured to be a part of Ethernet VDC on a Cisco Nexus 7000 Series device drop all FCoE traffic carried on that link. A network design using dedicated links for Ethernet and FCoE traffic makes use of the high-availability features of both LAN and SAN networks.

Because Cisco MDS 9000 Series devices do not support the switching of Ethernet traffic, we recommend that you ensure that no Ethernet VLANs get mapped to interfaces that connect Cisco Nexus 5000 Series and Cisco MDS 9000 Series devices. Directing Ethernet LAN traffic to a Cisco MDS 9000 Series devices that causes a traffic black hole and might result in loss of traffic.

Storage VDCs for Cisco Nexus 7000 Series Devices

VDCs that are available on Cisco Nexus 7000 Series devices allow virtualization at the device level, where logical entities are created to provide process separation and fault tolerance. For FCoE, dedicated links can be configured so that ingress traffic is processed in a separate and distinct VDC, which is the storage VDC.

On the Cisco Nexus 7000 Series device, you must create a storage VDC for FCoE traffic as follows:

```
N7K-1# configure terminal
N7K-1(config)# vdc fcoe_vdc type storage
Note: Creating VDC, one moment please...
N7K-1(config)# limit-resource module-type f1
N7K-1(config)# allow feature-set fcoe
N7K-1(config-vdc)# allocate interface ethernet 3/1-32
```

Moving ports causes all configuration that is associated to them in the source VDC to be removed.

FCoE and Enhanced vPC Considerations

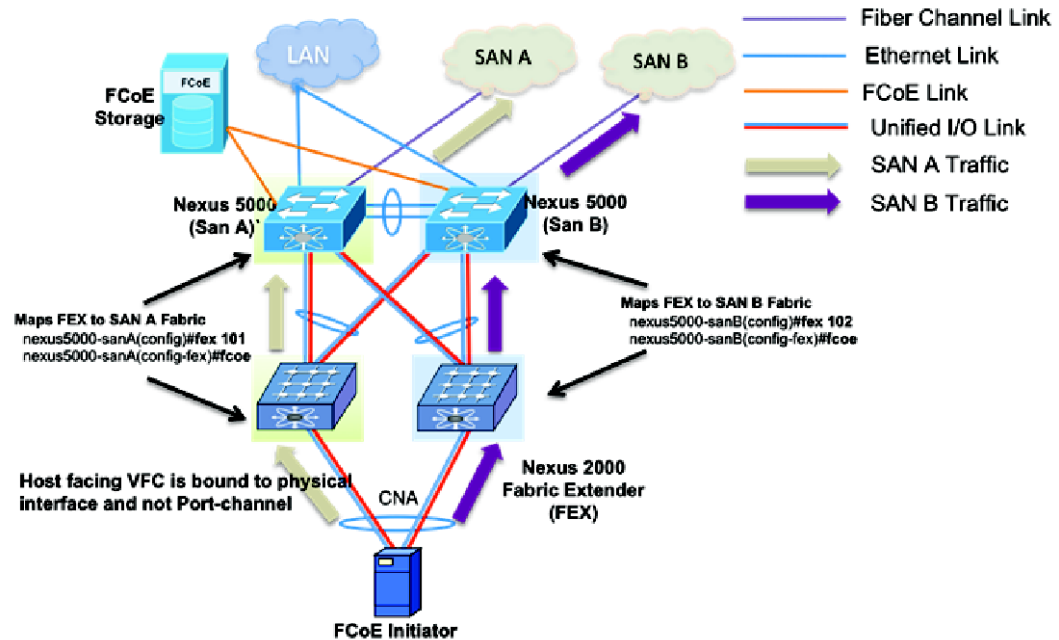
Beginning with the Cisco NX-OS Release 5.1(3)N1(1), Cisco Nexus 5000 Series devices support FCoE on Enhanced vPC (EVPC). With EVPC topologies, it is important to ensure that FCoE traffic does not get forwarded from the FEXs to more than one Cisco Nexus 5000 Series device, because it violates the SAN fabric isolation rule.

Send documentation comments to n5kdocfeedback@cisco.com

This behavior is implemented by associating each FEX with only one Cisco Nexus 5000 Series device. The Cisco Nexus 2000 and Cisco Nexus 5000 pair is part of only one FCoE SAN fabric.

Figure 6-3 shows an FCoE network design that uses an eVPC-layer vPC. The server FCoE traffic received on the FEX is sent only to the associated Cisco Nexus 5000 Series device. Using the `fcoe` command for the FEX configuration associates the FEX device with one Cisco Nexus 5000 Series device. One vPC link is used for FCoE traffic, which helps to implement the SAN A and SAN B isolation.

Figure 6-3 FCoE Network Design Using eVPC-layer vPC



331780

This example shows how to associate a FEX and a Cisco Nexus 5000 Series device with FCoE SAN fabrics:

```

nexus5000-sanA(config)# fex 101
nexus5000-sanA(config-fex)# fcoe
nexus5000-sanB(config)# fex 102
nexus5000-sanB(config-fex)# fcoe
  
```



Note

Do not associate the same FEX with multiple Cisco Nexus 5000 Series devices because it brings down host-facing VFC links that are associated with FEX and can cause FCoE traffic disruption.

SAN Boot of Initiators via LACP-Based Host vPC

Beginning with Cisco NX-OS Release 5.1(3)N1(1), Cisco Nexus 5000 Series devices support the SAN boot of initiators on Link Aggregation Control Protocol (LACP) based vPC. This limitation is specific to LACP-based port channels. The host-facing VFC interfaces are bound to port channel members instead of the port channel itself. This binding ensures that the host-side VFC comes up during a SAN boot as soon as the link on the CNA/ Host Bus Adapter (HBA) comes up, without relying on the LACP-based port channel to form first.


Send documentation comments to n5kdocfeedback@cisco.com

Figure 6-4 shows a network design with an initiator configured to boot from SAN via a vPC. VFC1 is bound to physical interface Eth101/1/1 that is part of port channel 1. VFC interfaces come up as soon as the channel-member state of host-facing vPC comes up, which helps the initiator to boot from SAN via the vPC.

This example shows how to configure the host-facing FEX:

```
nexus5000-sanA(config)#fex 101
nexus5000-sanA(config-fex)#fcoe
nexus5000-sanA(config)#interface vfc 1
nexus5000-sanA(config-if)#bind interface eth101/1/1
nexus5000-sanA(config)#interface eth101/1/1
nexus5000-sanA(config-if)#channel-group 1
nexus5000-sanB(config)#fex 102
nexus5000-sanB(config-fex)#fcoe
nexus5000-sanB(config)#interface vfc 1
nexus5000-sanB(config-if)#bind interface eth102/1/1
nexus5000-sanB(config)#interface eth102/1/1
nexus5000-sanB(config-if)#channel-group 1
```

Figure 6-4 Network Design with an Initiator Configured to Boot from SAN via a vPC



Note

VFC binding to host vPC configurations are permitted only if the FEXs are configured in a FEX straight-through topology (non-EVPC mode). This functionality ensures backward compatibility of configurations and all supported topologies prior to Cisco NX-OS release 5.1(3)N1(1).



Note

Binding multiple VFC interfaces with multiple members of a vPC is not allowed.

Send documentation comments to n5kdocfeedback@cisco.com



Note

VFC binding to a member of a port channel in a two-layer vPC topology is not allowed if the member port resides on a FEX that is not associated with the local Cisco Nexus 5000 Series device.

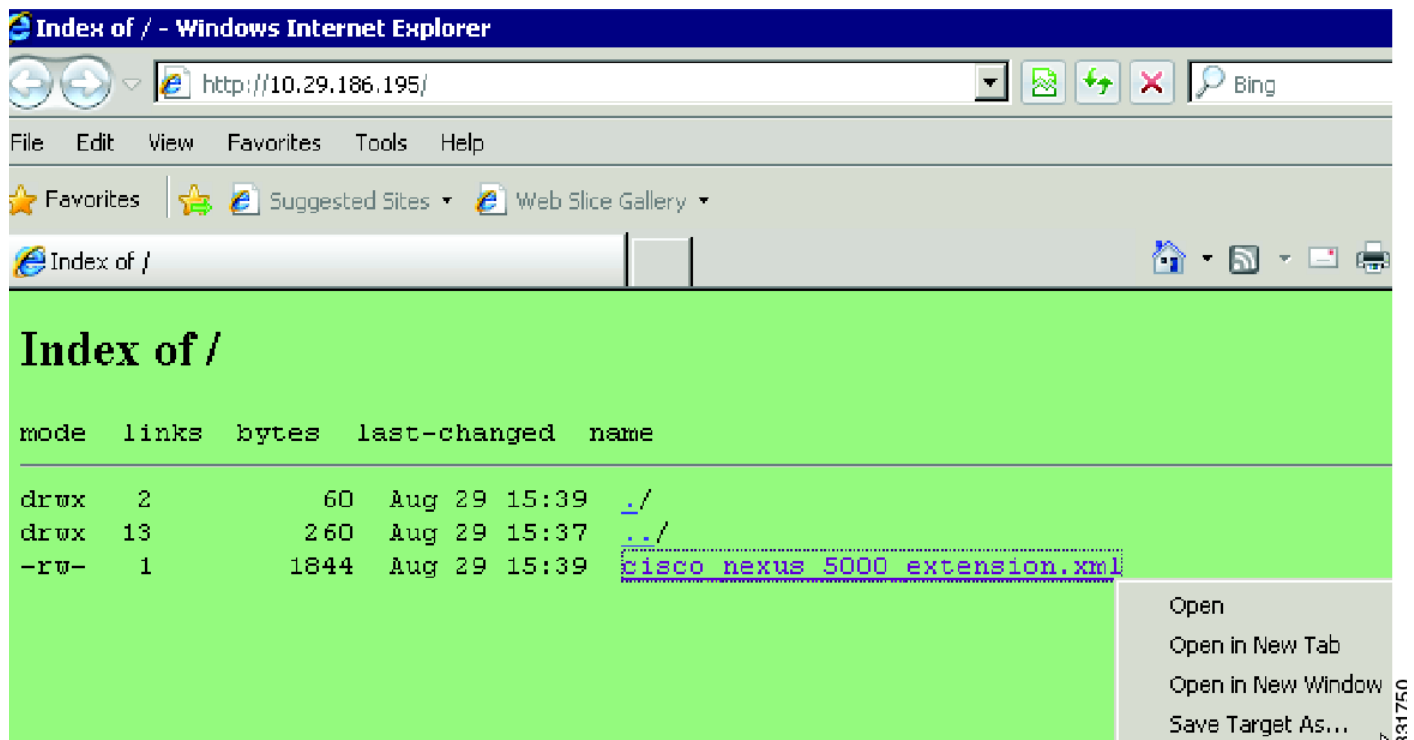
FCoE Functionality Using the Adapter-FEX

Introduction

Beginning with Cisco NXOS 5.1(3)N1(1), Cisco Nexus 5000 Series devices, when connected to Cisco UCS servers with P81E adapters, can support FCoE over vEthernet interfaces. The Cisco UCS P81E Virtual Interface Card is a virtualization-optimized Fibre Channel over Ethernet (FCoE) PCI Express (PCIe) 2.0 x8 10-Gbps adapter designed for use with Cisco UCS C-Series rack-mount servers. When configured in Adapter-FEX mode, the two unified ports can be sliced into two FCoE channels, which creates four vHBA interfaces available to the operating system.

Figure 6-5 shows a topology implementation of FCoE over the Adapter-FEX. For FCoE over Adapter-FEX network designs, the current supported configuration is to use a vPC between the Cisco Nexus Series 5000 Series and a FEX. At the same time, the host end of the P81E adapter should be configured in Active/Standby mode. For FCoE and EVPC topologies, the FEX gets associated with one Cisco Nexus 5000 Series fabric for FCoE traffic forwarding. Because the host HBA is configured in Active/Standby mode, a separate Ethernet to vEthernet to VFC mapping is needed on the other FEX and Cisco Nexus 5000 Series pair for topology implementation of FCoE over Adapter-FEX

Figure 6-5 Topology implementation of FCoE over Adapter-FEX



Send documentation comments to n5kdocfeedback@cisco.com

The following configuration is required on the Cisco Nexus 5000 Series device for binding the VFC to vEthernet interfaces. For FCoE, the Adapter-FEX can be deployed using the Active/Standby mode on the host side.

```
nexus5000-sanA(config)# install feature-set virtualization
nexus5000-sanA(config)# feature-set virtualization
nexus5000-sanA(config)# feature vmfex
nexus5000-sanA(config)# interface vfc 1
nexus5000-sanA(config-if)# bind interface veth 1
nexus5000-sanA(config-if)# interface veth 1
nexus5000-sanA(config-if)# bind interface eth101/1/1 channel 3
nexus5000-sanA(config-if)# interface eth101/1/1
nexus5000-sanA(config-if)# switchport mode vntag
nexus5000-sanA(config-if)# fex 101
nexus5000-sanA(config-fex)# fcoe

nexus5000-sanB(config)# interface vfc 2
nexus5000-sanB(config-if)# bind interface veth 2
nexus5000-sanB(config-if)# interface veth 2
nexus5000-sanB(config-if)# bind interface eth102/1/1 channel 4
nexus5000-sanB(config-if)# interface eth102/1/1
nexus5000-sanB(config-if)# switchport mode vntag
nexus5000-sanB(config-if)# fex 102
nexus5000-sanB(config-fex)# fcoe
```

**Note**

Because the vNIC of a virtual machine is mapped to the vEthernet interface on Cisco Nexus 5000 Series devices, between the two Cisco Nexus 5000 Series devices configured for redundancy, the vEthernet interface number has to be unique.

**Note**

FCoE and FCoE NPV functionalities can co exist with FabricPath by using a dedicated link for FCoE traffic.

**Note**

NPIV mode is not supported on the Cisco UCS P81E adapter when connected to a Cisco Nexus 5000 Series device using a VNP port.

**Note**

SAN boot of initiators using the Cisco UCS P81E virtual interface card is supported only when the Network Interface Virtualization feature is enabled.

Creating FCoE Port Profiles on Cisco Nexus 5000 Series Devices

The following configurations are required to create vEthernet port profiles on the Cisco Nexus 5000 Series devices and to form connectivity to vCenter:

Port Profile Configuration:

```
nexus5000-sanA(config)# config t
nexus5000-sanA(config)# port-profile type vethernet vnic-fcoe-1
nexus5000-sanA(config)# switchport mode trunk
nexus5000-sanA(config)# switchport trunk allowed vlan 1,100
nexus5000-sanA(config)# state enabled
```

SVS connectivity to vCenter:

Send documentation comments to n5kdocfeedback@cisco.com

```
nexus5000-sanA(config)# svcs connection vCenter-Nexus5000
nexus5000-sanA(config)# protocol vmware-vim
nexus5000-sanA(config)# remote ip address 172.28.3.19 port 80 vrf management
nexus5000-sanA(config)# dvs-name SJC-LAB
nexus5000-sanA(config)# vmware dvs datacenter-name TME-LAB
nexus5000-sanA(config)# connect
```



Note

Configuring SVS connectivity with a vCenter automatically pushes the port profiles created on the Cisco Nexus 5000 Series device to the vCenter.

Creating vHBAs on Cisco UCS Servers and Binding to Port Profiles

The Adapter-FEX feature of Cisco UCS P81E adapters on UCS servers is configured by using the Cisco UCS Manager CLI. This example shows how to connect to the Cisco UCS Manager CLI and to create vHBA interfaces:

```
sjc-xdm-054$ ssh -l admin ucs-afex-02.cisco.com
admin@ucs-afex-02.cisco.com's password:
ucs-afex-02# scope chassis
ucs-afex-02 /chassis # show adapter

PCI Slot Product Name Serial Number Product ID Vendor
-----
1 UCS VIC P81E QCI1532A34U N2XX-ACPCI01 Cisco Systems Inc
ucs-afex-02 /chassis # scope adapter 1
ucs-afex-02 /chassis/adapter # set niv-mode enabled
ucs-afex-02 /chassis/adapter # create host-fc-if vHBA-P81E-01
ucs-afex-02 /chassis/adapter # set port-profile vnic-fcoe-1
ucs-afex-02 /chassis/adapter/host-fc-if *# set channel-number 7
ucs-afex-02 /chassis/adapter/host-fc-if *# commit
ucs-afex-02 /chassis/adapter/host-fc-if # show detail
Name vHBA-P81E-01:
World Wide Node Name: 10:00:E8:B7:48:4E:0A:A0
World Wide Port Name: 20:00:E8:B7:48:4E:0A:A0
FC SAN Boot: disabled
Persistent LUN Binding: disabled
Uplink Port: 0
MAC Address: E8:B7:48:4E:0A:A0
CoS: N/A
VLAN: NONE
Rate Limiting: N/A
PCIe Device Order: ANY
EDTOV: 2000
RATOV: 10000
Maximum Data Field Size: 2048
Channel Number: 5
Port Profile: vnic-fcoe-1
```



Note

Creation of more than two vHBA interfaces is not supported using the Cisco UCS C-Series Cisco Integrated Management Controller (CIMC) graphical user interface (GUI) in Cisco NX-OS Release OS 5.1(3)N1(1).

Send documentation comments to n5kdocfeedback@cisco.com

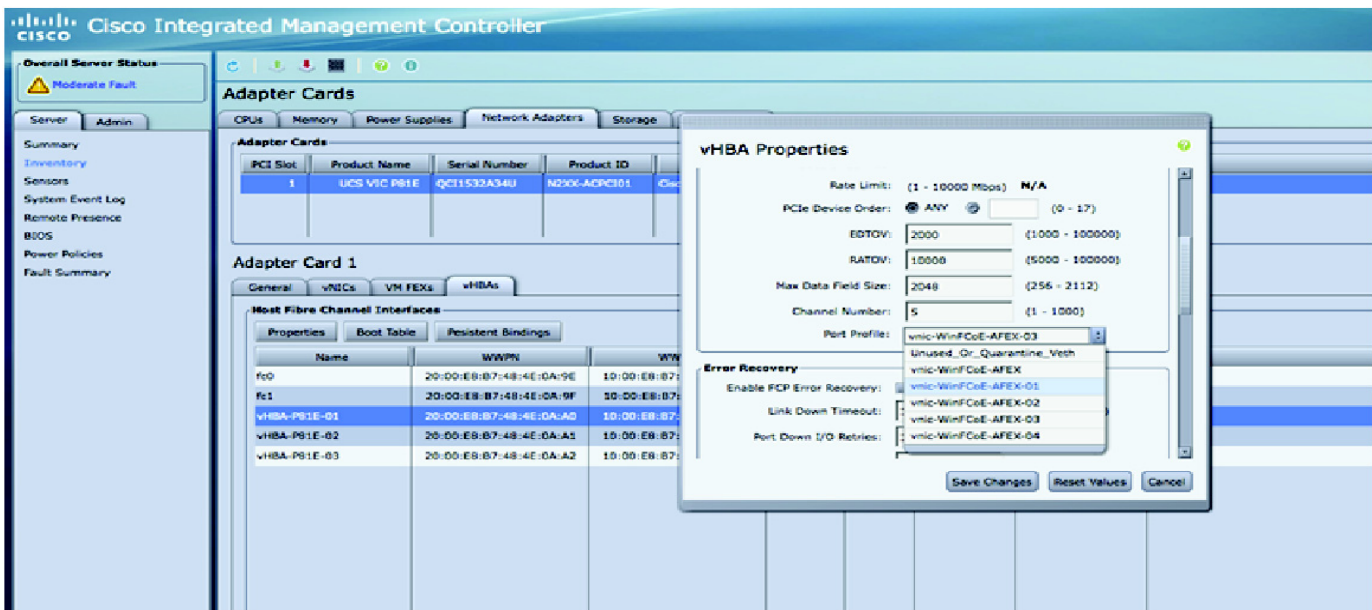
Binding vHBAs to Port Profiles Using CIMC

Virtual HBA interfaces created using the UCS C-Series Server CLI can be mapped to port profiles using the CIMC GUI. CIMC is the management service for the UCS C-Series server and runs within the server. The CIMC component can be accessed using a web browser.

Port profiles created on Cisco Nexus 5000 Series devices are pushed to the Cisco UCS P81E adapter using a VNTag link. VNTag is a unique tag that is assigned by the adapter to identify the source and destination vNIC. Port profiles appear in the drop-down menu of CIMC.

Figure 6-6 shows how to map port profiles to vHBA interfaces using the CIMC tool.

Figure 6-6 Mapping port profiles to vHBA interfaces using the CIMC tool



381783



INDEX

B

- buffer allocation
 - and QoS configuration [1-11](#)
 - configuring for FCoE COS [1-8](#)
 - for FCoE [1-8](#)

C

- class-fcoe [1-8](#)
- class of service (COS) [1-12](#)
 - and ETS [1-12](#)
 - and PFC [1-12](#)
- CNA
 - DCB support [1-13](#)
 - second generation [1-6](#)
- consolidated links [1-9, 1-10](#)
 - benefits [1-10](#)
- consolidated vs dedicated links [1-9](#)
- COS
 - default value and FCoE [1-12](#)

D

- Data Center Bridging eXchange (DCBX) [1-13](#)
- DCB Ethernet links [1-9](#)
- DCBX
 - negotiation failure [1-14](#)
- dedicated links [1-9, 1-10](#)
 - benefits [1-10](#)
- default mode
 - on Cisco Nexus 5000 Series switch [1-15](#)
- Domain IDs

- limitations [1-15](#)

E

- Enhanced Transmission Selection (ETS) [1-12](#)
- Ethernet NIC [1-6](#)
- ETS
 - default settings [1-12](#)

F

- FC-MAP [1-2](#)
 - changing the FC-MAP value [1-2](#)
 - default value [1-2](#)
 - ranges [1-2](#)
- FCoE
 - buffer allocation [1-8](#)
 - enabling [1-2](#)
 - enabling on VLAN 1 [1-3](#)
 - host disruptions [1-2](#)
 - interoperability [1-14](#)
 - no-drop class of service
 - and QoS configuration example [1-12](#)
 - predefined QoS policies [1-11](#)
 - QoS configuration [1-11](#)
 - single-hop topology [1-14](#)
- FCoE fabric
 - best practice [1-3](#)
 - configuring [1-3](#)
- FCoE ports
 - host-facing [1-3](#)
- FCoE VLAN
 - and STP [1-3, 1-4](#)

Send documentation comments to n5kdocfeedback@cisco.com

configuration in a vPC [1-7](#)

connecting to a VF port [1-3](#)

FCoE VLANs

difference from Ethernet VLANs [1-3](#)

Fibre Channel

HBA [1-6](#)

H

high availability (HA) [1-2](#)

I

IEEE 802.1Q Data Center Bridging (DCB) standard [1-13](#)

IEEE 802.1Q Enhance Ethernet Standards [1-12](#)

IEEE 802.1Q standard [1-12](#)

interoperability

and FCoE [1-14, 1-15](#)

L

link aggregation control protocol (LACP) [1-5](#)

load balance [1-2](#)

M

MST [1-4](#)

N

native

fabric services [1-15](#)

network disruptions [1-2](#)

new and changed features (table) [2-9](#)

no-drop classes of service [1-12](#)

no-drop service

thresholds [1-8](#)

N-Port ID Virtualization (NPIV) [1-15](#)

N-Port Virtualizer (NPV) [1-15](#)

NPV

device benefits [1-15](#)

NPIV requirement [1-15](#)

NPV mode

changing to switch mode [1-15](#)

requirement [1-15](#)

P

PFC

class-of-service [1-12](#)

default settings [1-12](#)

lossless transport and dedicated bandwidth [1-12](#)

pre-defined

FCoE policies [1-11](#)

Priority Flow Control (PFC) [1-12](#)

PVST [1-4](#)

PVST+ [1-4](#)

Q

QoS

FCoE configuration [1-11](#)

S

single-hop

FCoE topology [1-14](#)

Spanning Tree Protocol [1-3](#)

switch mode

and FCoE [1-15](#)

and native fabric services [1-15](#)

changing to NPV mode [1-15](#)

U

Unified Port Controller (UPC) ASIC [1-10](#)

Send documentation comments to n5kdocfeedback@cisco.com

first generation [1-10](#)

second generation [1-10](#)

VLAN configuration limit [1-11](#)

unified ports [1-11](#)

configuration requirements [1-11](#)

in expansion modules [1-11](#)

V

Virtual Port Channeling (vPC)

and FCoE [1-5](#)

VLAN

scalability [1-11](#)

VLAN to VSAN mapping [1-3](#)

vPC

connecting a host [1-5](#)

Send documentation comments to n5kdocfeedback@cisco.com